On Matrices With Displacement Structure: Generalized Operators and Faster Algorithms

Alin Bostan, Claude-Pierre Jeannerod, Christophe Moulleron, Eric Schost

► To cite this version:

Alin Bostan, Claude-Pierre Jeannerod, Christophe Moulleron, Eric Schost. On Matrices With Displacement Structure: Generalized Operators and Faster Algorithms. SIAM Journal on Matrix Analysis and Applications, 2017, 38 (3), pp.733-775. 10.1137/16M1062855 . hal-01588552

HAL Id: hal-01588552

<https://hal.science/hal-01588552>

Submitted on 15 Sep 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ON MATRICES WITH DISPLACEMENT STRUCTURE: GENERALIZED OPERATORS AND FASTER ALGORITHMS

A. BOSTAN*, C.-P. JEANNEROD†, C. MOUILLERON‡, AND É. SCHOST§

Abstract. For matrices with displacement structure, basic operations like multiplication, inversion, and linear system solving can all be expressed in terms of the following task: evaluate the product AB , where A is a structured $n \times n$ matrix of displacement rank α , and B is an arbitrary $n \times \alpha$ matrix. Given B and a so-called *generator* of A , this product is classically computed with a cost ranging from $O(\alpha^2 \mathcal{M}(n))$ to $O(\alpha^2 \mathcal{M}(n) \log(n))$ arithmetic operations, depending on the type of structure of A ; here, $\mathcal{M}$ is a cost function for polynomial multiplication. In this paper, we first generalize classical displacement operators, based on block diagonal matrices with companion diagonal blocks, and then design fast algorithms to perform the task above for this extended class of structured matrices. The cost of these algorithms ranges from $O(\alpha^{\omega-1} \mathcal{M}(n))$ to $O(\alpha^{\omega-1} \mathcal{M}(n) \log(n))$, with ω such that two $n \times n$ matrices over a field can be multiplied using $O(n^\omega)$ field operations. By combining this result with classical randomized regularization techniques, we obtain faster Las Vegas algorithms for structured inversion and linear system solving.

Key words. structured linear algebra, matrix multiplication, computational complexity

AMS subject classifications. 65F05, 68Q25

1. Introduction. Exploiting the structure of data is key to develop fast algorithms and, in the context of linear algebra, this principle is at the heart of algorithms for *displacement structured matrices*. These algorithms can speed up for instance the inversion of a given matrix whenever this matrix has a structure close to that of a Toeplitz, Hankel, Vandermonde, or Cauchy matrix. The idea is to represent structured matrices succinctly by means of their *generators* with respect to suitable *displacement operators*, and to operate on this succinct data structure.

Displacement operators. Let $\mathbb{F}$ be a field. To measure the extent to which a matrix $A \in \mathbb{F}^{m \times n}$ possesses some structure, it is customary to use its *displacement rank*, that is, the rank of its image through a *displacement operator* [18]. There exist two broad classes of displacement operators: *Sylvester operators*, of the form

$$\nabla_{M,N} : A \in \mathbb{F}^{m \times n} \mapsto MA - AN \in \mathbb{F}^{m \times n},$$

and *Stein operators*, of the form

$$\Delta_{M,N} : A \in \mathbb{F}^{m \times n} \mapsto A - MAN \in \mathbb{F}^{m \times n};$$

in both cases, M and N are fixed matrices in $\mathbb{F}^{m \times m}$ and $\mathbb{F}^{n \times n}$, respectively. For any such operator, say $\mathcal{L}$, the rank of $\mathcal{L}(A)$ is called the $\mathcal{L}$ -*displacement rank* of A , or simply its displacement rank if $\mathcal{L}$ is clear from the context. Loosely speaking, the matrix A is called *structured* (with respect to the operator $\mathcal{L}$) if its $\mathcal{L}$ -displacement rank is small compared to its sizes m and n .

We say that a matrix pair (G, H) in $\mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha}$ is an $\mathcal{L}$ -*generator of length α* of $A \in \mathbb{F}^{m \times n}$ if it satisfies $\mathcal{L}(A) = GH^t$, with H^t the transpose of H . Again, if $\mathcal{L}$ is clear

*Inria, France (alin.bostan@inria.fr).

†Inria, Université de Lyon, laboratoire LIP (CNRS, ENSL, Inria, UCBL), France (claude-pierre.jeannerod@inria.fr).

‡ENSIIE, Évry, France (christophe.mouilleron@ens-lyon.org). Part of this work was done when C. Mouilleron was a member of laboratoire LIP (CNRS, ENSL, Inria, UCBL).

§David R. Cheriton School of Computer Science, University of Waterloo, ON, Canada (eschost@uwaterloo.ca).

from the context, we shall simply say *generator of length α* . Provided $\mathcal{L}$ is invertible and when α is small, such a generator can play the role of a succinct data structure to represent the matrix $\mathbf{A}$. The smallest possible value for α is the rank of $\mathcal{L}(\mathbf{A})$.

If $\mathbf{A}$ is invertible and structured with respect to $\mathcal{L}$, then its inverse $\mathbf{A}^{-1}$ is structured with respect to the operator $\mathcal{L}'$ obtained by swapping $\mathbf{M}$ and $\mathbf{N}$:

$$(1) \quad \mathcal{L} = \nabla_{\mathbf{M}, \mathbf{N}} \implies \mathcal{L}' = \nabla_{\mathbf{N}, \mathbf{M}}, \quad \mathcal{L} = \Delta_{\mathbf{M}, \mathbf{N}} \implies \mathcal{L}' = \Delta_{\mathbf{N}, \mathbf{M}}.$$

More precisely, the above claim says that the ranks of $\mathcal{L}(\mathbf{A})$ and $\mathcal{L}'(\mathbf{A}^{-1})$ are the same. (See [33, Theorem 1.5.3] and, for the case not handled there, see [18, p. 771]; for completeness, we give a proof in Appendix A.)

Some classical operators. Consider first Toeplitz-like matrices. For φ in $\mathbb{F}$, it is customary to define the $m \times m$ cyclic down-shift matrix

$$(2) \quad \mathbb{Z}_{m, \varphi} = \begin{bmatrix} & & & \varphi \\ 1 & & & \\ & \ddots & & \\ & & 1 & \end{bmatrix} \in \mathbb{F}^{m \times m}.$$

Then, using Sylvester operators, a matrix $\mathbf{A} \in \mathbb{F}^{m \times n}$ will be called Toeplitz-like if $\mathbb{Z}_{m, \varphi} \mathbf{A} - \mathbf{A} \mathbb{Z}_{n, \psi}$ has a low rank compared to m and n . (This rank is independent of the choice of φ and ψ , up to an additive constant of absolute value at most two.) Using Stein operators, one would consider instead the rank of $\mathbf{A} - \mathbb{Z}_{m, \varphi} \mathbf{A} \mathbb{Z}_{n, \psi}^t$.

Besides $\mathbb{Z}_{m, \varphi}$ and its transpose it is also useful to consider the diagonal matrix $\mathbb{D}(\mathbf{x})$, whose diagonal is given by $\mathbf{x} \in \mathbb{F}^m$. Specifically, popular choices take

$$(3) \quad (\mathbf{M}, \mathbf{N}) \in \{\mathbb{Z}_{m, \varphi}, \mathbb{Z}_{m, \varphi}^t, \mathbb{D}(\mathbf{x})\} \times \{\mathbb{Z}_{n, \psi}, \mathbb{Z}_{n, \psi}^t, \mathbb{D}(\mathbf{y})\},$$

with $\varphi, \psi \in \mathbb{F}$, $\mathbf{x} \in \mathbb{F}^m$ and $\mathbf{y} \in \mathbb{F}^n$, for either Sylvester or Stein operators. This family covers in particular Toeplitz and Hankel-like structures, where both $\mathbf{M}$ and $\mathbf{N}$ are (transposed) cyclic down-shift matrices; Vandermonde-like structures, where one of the matrices $\mathbf{M}$ and $\mathbf{N}$ is a (transposed) cyclic down-shift matrix and the other one is diagonal; and Cauchy-like structures, where both $\mathbf{M}$ and $\mathbf{N}$ are diagonal. We say that a displacement operator is of *Toeplitz / Hankel type* if it falls into the first category, and of *Vandermonde / Cauchy type* if it falls into the second or third one.

Block-diagonal companion matrices. The first contribution of this paper is to generalize the classical displacement operators seen before. Let $\mathbb{F}[x]$ be the univariate polynomial ring over $\mathbb{F}$, and for $\delta \geq 0$, let $\mathbb{F}[x]_\delta$ be the $\mathbb{F}$ -vector space of polynomials of degree less than δ . Furthermore, for $F = \sum_i f_i x^i \in \mathbb{F}[x]$ monic of degree $\delta > 0$, denote by $\mathbb{M}_F$ the $\delta \times \delta$ companion matrix associated with F :

$$\mathbb{M}_F = \begin{bmatrix} & & -f_0 \\ & & -f_1 \\ & \ddots & \vdots \\ & & 1 & -f_{\delta-1} \end{bmatrix} \in \mathbb{F}^{\delta \times \delta}.$$

Two special cases are the $m \times m$ matrix $\mathbb{Z}_{m, \varphi}$ defined in (2) and the 1×1 matrix given by $x_0 \in \mathbb{F}$, which are associated with the polynomials $x^m - \varphi$ and $x - x_0$, respectively.

Given $d > 0$, let now $\mathbf{P} = P_1, \dots, P_d$ denote a family of monic nonconstant polynomials in $\mathbb{F}[x]$, and let $m = m_1 + \dots + m_d$ with $m_i = \deg(P_i) \geq 1$ for all i . We

will call *block-diagonal companion matrix* associated with $\mathbf{P}$ the $m \times m$ block-diagonal matrix $\mathbb{M}_{\mathbf{P}}$ whose i th diagonal block is the $m_i \times m_i$ companion matrix $\mathbb{M}_{P_i}$:

$$\mathbb{M}_{\mathbf{P}} = \begin{bmatrix} \mathbb{M}_{P_1} & & \\ & \ddots & \\ & & \mathbb{M}_{P_d} \end{bmatrix} \in \mathbb{F}^{m \times m}.$$

We write P to denote the product $P_1 \cdots P_d$, which is the characteristic polynomial of $\mathbb{M}_{\mathbf{P}}$. Finally, we associate with $\mathbf{P}$ the following assumption on its elements:

$$\mathcal{H}_{\mathbf{P}}: \quad P_1, \dots, P_d \text{ are pairwise coprime.}$$

Associated displacement operators. With $\mathbf{P}$ as above, we now consider another family of monic nonconstant polynomials, namely $\mathbf{Q} = Q_1, \dots, Q_e$, with respective degrees $n_1, \dots, n_e$; we let $Q = Q_1 \cdots Q_e$ and $n = n_1 + \cdots + n_e$. In what follows, we assume that both assumptions $\mathcal{H}_{\mathbf{P}}$ and $\mathcal{H}_{\mathbf{Q}}$ hold.

Let then $\mathbb{M}_{\mathbf{P}} \in \mathbb{F}^{m \times m}$ and $\mathbb{M}_{\mathbf{Q}} \in \mathbb{F}^{n \times n}$ be the block-diagonal companion matrices associated with $\mathbf{P}$ and $\mathbf{Q}$. In this paper, we consider the following eight operators:

$$\begin{aligned} & \nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}, \quad \nabla_{\mathbb{M}_{\mathbf{P}}^t, \mathbb{M}_{\mathbf{Q}}}, \quad \nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}}, \quad \nabla_{\mathbb{M}_{\mathbf{P}}^t, \mathbb{M}_{\mathbf{Q}}^t}, \\ & \Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}, \quad \Delta_{\mathbb{M}_{\mathbf{P}}^t, \mathbb{M}_{\mathbf{Q}}}, \quad \Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}}, \quad \Delta_{\mathbb{M}_{\mathbf{P}}^t, \mathbb{M}_{\mathbf{Q}}^t}. \end{aligned}$$

We shall call them the *operators associated with $(\mathbf{P}, \mathbf{Q})$ of Sylvester or Stein type*, respectively, and say that they have *format (m, n)* . Two of these operators will be highlighted, $\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ and $\Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$, and we will call them the *basic operators* associated with $(\mathbf{P}, \mathbf{Q})$. Indeed, we will be able to derive convenient formulas to invert them; this will in turn allow us to deal with the other six operators by suitable reductions.

Tables 1 and 2 show that the classical structures defined by (3) follow as special cases of these operators, for suitable choices of $(\mathbf{P}, \mathbf{Q})$ making $\mathbb{M}_{\mathbf{P}}$ or $\mathbb{M}_{\mathbf{Q}}$ a diagonal or a (transposed) cyclic down-shift matrix.

TABLE 1
Some particular cases for the Sylvester operator $\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$.

	$e = 1$ and $Q_1 = x^n - \psi$	$e = n$ and $Q_j = x - y_j$
$d = 1$ and $P_1 = x^m - \varphi$	Hankel-like	Vandermonde ^t -like (for the points $1/y_j$)
$d = m$ and $P_i = x - x_i$	Vandermonde-like (for the points $1/x_i$)	Cauchy-like (for the points x_i and y_j)

TABLE 2
Some particular cases for the Stein operator $\Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$.

	$e = 1$ and $Q_1 = x^n - \psi$	$e = n$ and $Q_j = x - y_j$
$d = 1$ and $P_1 = x^m - \varphi$	Toeplitz-like	Vandermonde ^t -like (for the points y_j)
$d = m$ and $P_i = x - x_i$	Vandermonde-like (for the points x_i)	Cauchy-like (for the points $1/x_i$ and y_j)

In other words, our family of operators covers all classical operators of Toeplitz, Hankel, Vandermonde, and Cauchy type in a unified manner. However, this formalism also includes many new cases as “intermediate” situations. For example, taking $e = 1$ and $Q_1 = x^n$, we see that the matrix of the *multiple reduction map* modulo $P_1, \dots, P_d$ has displacement rank 1 for the Stein operator $\Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$, so our operators will allow us to address problems related to Chinese remaindering.

Problem statement. Our goal is to study the complexity of basic computations with matrices that are structured for the operators seen above: multiply a matrix $A \in \mathbb{F}^{m \times n}$ by a matrix B , invert A , and solve the linear system $Ax = b$. Formally, let $\mathcal{L} : \mathbb{F}^{m \times n} \rightarrow \mathbb{F}^{m \times n}$ be a displacement operator of either Sylvester or Stein type, and assume that $\mathcal{L}$ is invertible. Then, the problems we address read as follows.

$\text{mul}(\mathcal{L}, \alpha, \beta)$: given an $\mathcal{L}$ -generator $(G, H) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha}$ of a matrix $A \in \mathbb{F}^{m \times n}$ with $\alpha \leq \min(m, n)$ and given a matrix $B \in \mathbb{F}^{n \times \beta}$, compute the product $AB \in \mathbb{F}^{m \times \beta}$.

$\text{inv}(\mathcal{L}, \alpha)$: given an $\mathcal{L}$ -generator $(G, H) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{m \times \alpha}$ of a matrix $A \in \mathbb{F}^{m \times m}$ with $\alpha \leq m$, return an $\mathcal{L}'$ -generator (G', H') of length at most α for A^{-1} , or assert that A is not invertible. Here, $\mathcal{L}'$ is the operator in (1).

$\text{solve}(\mathcal{L}, \alpha)$: given an $\mathcal{L}$ -generator $(G, H) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha}$ of a matrix $A \in \mathbb{F}^{m \times n}$ with $\alpha \leq \min(m, n)$ and given a vector $b \in \mathbb{F}^m$, find a solution $x \in \mathbb{F}^n$ to $Ax = b$, or assert that no such solution exists. (When $b = 0$ and A has not full column rank, then a nonzero vector x should be returned; we call this a *nontrivial solution*.)

For the problems above to make sense, we need the displacement operator $\mathcal{L}$ to be invertible. For the Sylvester operators associated with $(\mathbf{P}, \mathbf{Q})$, this occurs if and only if $\gcd(P, Q) = 1$; for Stein operators, the condition becomes $\gcd(P, \tilde{Q}) = 1$ with $\tilde{Q} = x^n Q(1/x)$ or, equivalently, $\gcd(\tilde{P}, Q) = 1$ with $\tilde{P} = x^m P(1/x)$ [33, Theorem 4.3.2].

We allow probabilistic algorithms. For instance, for inversion in size m , we will see algorithms that use $r = O(m)$ random elements in $\mathbb{F}$, for which success is conditional to avoiding a hypersurface in $\overline{\mathbb{F}}^r$ of degree $m^{O(1)}$.

To analyze all upcoming algorithms, we count all arithmetic operations in $\mathbb{F}$ at unit cost, so the *time* spent by an algorithm is simply the number of such operations it performs. Then, our goal is to give upper bounds on the cost functions $\mathcal{C}_{\text{mul}}(\cdot, \cdot, \cdot)$, $\mathcal{C}_{\text{solve}}(\cdot, \cdot)$, $\mathcal{C}_{\text{inv}}(\cdot, \cdot)$, which are such that the problems $\text{mul}(\mathcal{L}, \alpha, \beta)$, $\text{solve}(\mathcal{L}, \alpha)$, $\text{inv}(\mathcal{L}, \alpha)$, can be solved in respective times

$$\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta), \quad \mathcal{C}_{\text{solve}}(\mathcal{L}, \alpha), \quad \mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha).$$

In particular, $\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, 1)$ denotes the cost of a structured matrix-vector product Av .

A few further problems can be solved as direct extensions of these operations. Indeed, a simple transformation (see e.g. [33, Theorem 1.5.2]) shows that if a matrix A is structured with respect to one of the operators associated with $(\mathbf{P}, \mathbf{Q})$, its transpose A^t is structured as well, with respect to one of the operators associated with $(\mathbf{Q}, \mathbf{P})$. This is summarized as follows, where $(M, N) \in \{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{P}}^t\} \times \{\mathbb{M}_{\mathbf{Q}}, \mathbb{M}_{\mathbf{Q}}^t\}$:

$$\begin{aligned} \nabla_{M, N}(A) = GH^t &\iff \nabla_{N^t, M^t}(A^t) = (-H)G^t, \\ \Delta_{M, N}(A) = GH^t &\iff \Delta_{N^t, M^t}(A^t) = HG^t. \end{aligned}$$

Thus, from our results on multiplication, inversion, and linear system solving for the matrix A , one directly obtains similar results for the same operations with A^t .

Cost measures. We let $\mathcal{M}$ be a multiplication time function for $\mathbb{F}[x]$, that is, $\mathcal{M} : \mathbb{N}_{>0} \rightarrow \mathbb{R}_{>0}$ is such that two polynomials of degree less than d can be multiplied in $\mathcal{M}(d)$ arithmetic operations in $\mathbb{F}$; $\mathcal{M}$ must also satisfy the super-linearity

properties of [12, Ch. 8]. It follows from [38, 37] that we can always take $\mathcal{M}(d) \in O(d \lg(d) \lg \lg(d))$, where $\lg(d) = \log(\max(d, 2))$ and $\log$ is the binary logarithm. (This definition of $\lg(d)$ ensures that expressions like $d \lg(d) \lg \lg(d)$ do not vanish at $d = 1$.) Hence, from now on we assume that $\mathcal{M}$ is quasi-linear in d , so that $\mathcal{M}(d) = O(d^{1+\epsilon})$ for all $\epsilon > 0$, and that $\mathcal{M}(1) = 1$.

Here and hereafter, ω denotes any real number such that two $n \times n$ matrices over $\mathbb{F}$ can be multiplied using $O(n^\omega)$ arithmetic operations in $\mathbb{F}$. We have $\omega < 2.38$ [23] as $n \rightarrow \infty$ and, for moderate dimensions (say, $n < 10^6$), upper bounds on ω range from 2.81 [40] down to 2.78 [29]; on the other hand, we have the trivial lower bound $\omega \geq 2$. (Our cost analyses will be given for $\omega > 2$ for simplicity, but note that their extension to the case $\omega = 2$ would be straightforward and imply no more than some extra logarithmic factors.)

The running time of our algorithms will depend on the cost of some polynomial matrix operations. Let $\mathcal{M}_{\text{mat}}(d, n) : \mathbb{R}_{\geq 1} \times \mathbb{N}_{>0} \rightarrow \mathbb{R}_{>0}$ be such that two $n \times n$ matrices over $\mathbb{F}[x]_d$ can be multiplied in $\mathcal{M}_{\text{mat}}(d, n)$ arithmetic operations in $\mathbb{F}$. (It will be convenient to allow non-integer values of d ; this does not change the definition.) One can use the following estimates:

$$\mathcal{M}_{\text{mat}}(d, n) = \begin{cases} O(\mathcal{M}(d)n^\omega) & \text{in general,} \\ O(dn^\omega + \mathcal{M}(d)n^2) & \text{if } \text{char}(\mathbb{F}) = 0 \text{ or } |\mathbb{F}| \geq 2d; \end{cases}$$

the former follows from [9] and the latter is from [7], using evaluation and interpolation at geometric sequences; in both cases, this is $\tilde{O}(dn^\omega)$, where the soft-O notation $\tilde{O}(\cdot)$ means that we omit polylogarithmic factors. Our costs will be expressed using not quite $\mathcal{M}_{\text{mat}}$, but two related functions $\mathcal{M}'_{\text{mat}}$ and $\mathcal{M}''_{\text{mat}}$, which should be thought of as being “close” to $\mathcal{M}_{\text{mat}}$ up to logarithmic factors. These two functions are defined as follows. For n a power of two,

$$\mathcal{M}'_{\text{mat}}(d, n) = \sum_{k=0}^{\log(n)} 2^k \mathcal{M}_{\text{mat}}\left(2^k d, \frac{n}{2^k}\right)$$

and, for general n , $\mathcal{M}'_{\text{mat}}(d, n) = \mathcal{M}'_{\text{mat}}(d, \bar{n})$ with $\bar{n}$ the smallest power of two greater than or equal to n . Using the super-linearity of $\mathcal{M}$, the above choices for $\mathcal{M}_{\text{mat}}$ give

$$(4) \quad \mathcal{M}'_{\text{mat}}(d, n) = \begin{cases} O(\mathcal{M}(dn)n^{\omega-1}) & \text{in general,} \\ O(dn^\omega + \mathcal{M}(dn)n \lg(n)) & \text{if } \text{char}(\mathbb{F}) = 0 \text{ or } |\mathbb{F}| \geq 2dn; \end{cases}$$

in both cases, this is again $\tilde{O}(dn^\omega)$. Note on the other hand that $dn^2 \leq n \mathcal{M}(dn) \leq \mathcal{M}'_{\text{mat}}(d, n)$. Now, for d a power of two, let

$$\mathcal{M}''_{\text{mat}}(d, n) = \sum_{k=0}^{\log(d)} 2^k \mathcal{M}'_{\text{mat}}\left(\frac{d}{2^k}, n\right)$$

and, for general d , let $\mathcal{M}''_{\text{mat}}(d, n) = \mathcal{M}''_{\text{mat}}(\bar{d}, n)$ with $\bar{d}$ the smallest power of two greater than or equal to d . Using the two previous bounds on $\mathcal{M}'_{\text{mat}}$, we deduce that

$$\mathcal{M}''_{\text{mat}}(d, n) = \begin{cases} O(\mathcal{M}(dn)n^{\omega-1} \lg(d)) & \text{in general,} \\ O((dn^\omega + \mathcal{M}(dn)n \lg(n)) \lg(d)) & \text{if } \text{char}(\mathbb{F}) = 0 \text{ or } |\mathbb{F}| \geq 2dn; \end{cases}$$

again, in both cases this is $\tilde{O}(dn^\omega)$. Conversely, we have $dn^\omega = O(\mathcal{M}_{\text{mat}}''(d, n))$, since $\mathcal{M}_{\text{mat}}''(d, n) \geq d \mathcal{M}_{\text{mat}}'(1, n) \geq d \mathcal{M}_{\text{mat}}(1, n)$. Remark also that if we assume, similarly to the super-linearity assumptions on $\mathcal{M}$, that for every positive integer n the function $d \mapsto \mathcal{M}_{\text{mat}}'(d, n)/d$ is nondecreasing, then we have in all cases the simple bound

$$\mathcal{M}_{\text{mat}}''(d, n) \leq \mathcal{M}_{\text{mat}}'(d, n)(1 + \log(d)).$$

Finally, note that for $d = 1$ and $\omega > 2$,

$$\mathcal{M}_{\text{mat}}''(1, n) = \mathcal{M}_{\text{mat}}'(1, n) = O(n^\omega).$$

To check the second equality, one may fix some ϵ such that $0 < \epsilon < \omega - 2$, so that $\mathcal{M}(m) = O(m^{1+\epsilon})$ and thus $\mathcal{M}_{\text{mat}}'(1, n) = O(\sum_{k \geq 0} n^\omega (2^k)^{2+\epsilon-\omega})$ with $2 + \epsilon - \omega < 0$.

To any family of polynomials $\mathbf{P} = P_1, \dots, P_d$ with $m_i = \deg(P_i) > 0$ for all i , we will associate two cost measures, $\mathcal{C}(\mathbf{P})$ and $\mathcal{D}(\mathbf{P})$, related to Chinese remaindering and (extended) GCDs; both range between $O(m)$ or $O(\mathcal{M}(m))$ and $O(\mathcal{M}(m)\lg(m))$, where we write $m = m_1 + \dots + m_d$ as before. In both cases, we give a general formula and point out particular cases with smaller estimates.

First, in the particular case where $P_i = x - x_i$ for all i with the x_i in geometric progression, we take $\mathcal{C}(\mathbf{P}) = \mathcal{M}(m)$. In all other cases, we take

$$\mathcal{C}(\mathbf{P}) = \mathcal{M}(m) \left(1 + \sum_{1 \leq i \leq d} \frac{m_i}{m} \log \left(\frac{m}{m_i} \right) \right).$$

Since the value of the above sum ranges from zero to $\log(d)$ (see for example [8, p. 38]), the function $\mathcal{C}(\mathbf{P})$ always satisfies

$$\mathcal{M}(m) \leq \mathcal{C}(\mathbf{P}) \leq \mathcal{M}(m)(1 + \log(d));$$

in particular, we have $\mathcal{C}(\mathbf{P}) = O(\mathcal{M}(m)\lg(d))$.

The definition of the function $\mathcal{D}(\mathbf{P})$ starts with a particular case as well: when $d = 1$ and $P_1 = x^m - \varphi$ with $\varphi \in \mathbb{F}$, we take $\mathcal{D}(\mathbf{P}) = m$. Otherwise, we write

$$\mathcal{D}(\mathbf{P}) = \sum_{1 \leq i \leq d} \mathcal{M}(m_i) \lg(m_i).$$

Table 3 displays the values of $\mathcal{C}(\mathbf{P})$ and $\mathcal{D}(\mathbf{P})$ in the two special cases mentioned above, but also their estimates when $d = O(1)$ or $m_i = O(1)$ for $1 \leq i \leq d$. We see that when $d = O(1)$ we are essentially in the best case for $\mathcal{C}(\mathbf{P})$ and in the worst case for $\mathcal{D}(\mathbf{P})$, and that the situation is reversed when $m_i = O(1)$ for all i . In other words, the families with best and worst cases for $\mathcal{D}(\mathbf{P})$ are the opposite of those for $\mathcal{C}(\mathbf{P})$.

Finally, when considering two families of polynomials $\mathbf{P}$ and $\mathbf{Q}$, we will write

$$\mathcal{C}(\mathbf{P}, \mathbf{Q}) = \mathcal{C}(\mathbf{P}) + \mathcal{C}(\mathbf{Q}) \quad \text{and} \quad \mathcal{D}(\mathbf{P}, \mathbf{Q}) = \mathcal{D}(\mathbf{P}) + \mathcal{D}(\mathbf{Q}).$$

Background work. Following landmark publications such as [24, 18, 25, 3], the literature on structured linear systems has vastly developed, and we refer especially to [33, 34] for comprehensive overviews. It was recognized early that a key ingredient for the development of fast algorithms was the ability to perform matrix-vector products efficiently. For Toeplitz-like and Hankel-like systems, this was done in [25, 3];

TABLE 3
Upper bounds for $\mathcal{C}(\mathbf{P})$ and $\mathcal{D}(\mathbf{P})$ in some special cases.

	$\mathcal{C}(\mathbf{P})$	$\mathcal{D}(\mathbf{P})$
$P_i = x - x_i$ for all i , with the x_i in geometric progression	$\mathcal{M}(m)$	m
$d = 1$ and $P_1 = x^m - \varphi$	$\mathcal{M}(m)$	m
$m_i = O(1)$ for all i	$O(\mathcal{M}(m)\lg(m))$	$O(m)$
$d = O(1)$	$O(\mathcal{M}(m))$	$O(\mathcal{M}(m)\lg(m))$

for Vandermonde and Cauchy structures, this is in [13, 14]. In our notation, these references establish bounds of the form

$$\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, 1) = O(\alpha \mathcal{M}(m)) \quad \text{or} \quad \mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, 1) = O(\alpha \mathcal{M}(m)\lg(m)),$$

for operators of format (m, m) of Toeplitz / Hankel type or Vandermonde / Cauchy type, respectively.

If a matrix $\mathbf{B}$ has β columns, then the product $\mathbf{AB}$ can be computed by performing β matrix-vector products with $\mathbf{A}$. In other words, we can take

$$\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) \leq \beta \mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, 1),$$

which yields the bounds $O(\alpha\beta\mathcal{M}(m))$ and $O(\alpha\beta\mathcal{M}(m)\lg(m))$ for the cases seen above.

Such matrix-matrix products, with $\beta \simeq \alpha$, are the main ingredients in the Morf / Bitmead-Anderson (MBA) algorithm to invert structured matrices [3, 25] of Toeplitz / Hankel type, which combines the displacement rank approach of [11, 18, 17] with Strassen's divide-and-conquer approach for dense matrix inversion [40]. This algorithm initially required several genericity conditions to hold but, based on the randomized regularization technique of Kaltofen and Saunders [21], it was then further extended in [19, 20] to handle arbitrary matrices (see also [2, pp. 204–208] and [33, §5.5–5.7]). For operators of format (m, m) of Toeplitz, Hankel, Vandermonde or Cauchy type, the cost of inversion $\mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha)$ can then (roughly speaking) be taken in

$$O(\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \alpha)\lg(m));$$

see [36] for the case of Vandermonde-like and Cauchy-like matrices, and [26, 31, 32] for a unified treatment of all four structures. Using the above upper bound on the cost of multiplication $\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \alpha)$, we see that this is $\alpha\lg(m)$ times the cost of a matrix-vector product (strictly speaking, some precautions are necessary to make such a statement; for instance, the algorithm becomes probabilistic).

This results in running times of the form $O(\alpha^2\mathcal{M}(m)\lg(m))$ or $O(\alpha^2\mathcal{M}(m)\lg(m)^2)$ for inverting respectively Toeplitz / Hankel-like matrices and Vandermonde / Cauchy-like matrices of size m and displacement rank α . In the Vandermonde / Cauchy case, another approach due to Pan [30] proceeds by reduction to the Toeplitz / Hankel case; the running time is then reduced to $O(\alpha^2\mathcal{M}(m)\lg(m))$.

Going beyond the four basic structures, Olshevsky and Shokrollahi introduced a common generalization thereof, where the displacement matrices take the form of block-diagonal Jordan matrices [28]. We will discuss their result in more detail after

stating our main theorems. For the moment, we mention that their paper gives an algorithm for matrix-vector product with running times ranging from

$$\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, 1) = O(\alpha \mathcal{M}(m)) \quad \text{to} \quad \mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, 1) = O(\alpha \mathcal{M}(m) \lg(m)),$$

depending on the block configuration of the supporting Jordan matrices; for the four basic structures seen above, we recover the running times seen before. In [27], this result is used to sketch an extension of the MBA algorithm to such matrices. As before, the running time for inversion increases by a factor $\alpha \lg(m)$ compared to the time for matrix-vector multiplication, ranging from

$$\mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha) = O(\alpha^2 \mathcal{M}(m) \lg(m)) \quad \text{to} \quad \mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha) = O(\alpha^2 \mathcal{M}(m) \lg(m)^2),$$

depending on the operator.

When $\alpha = O(1)$, all above results for matrix-vector product or matrix inversion are within a polylogarithmic factor of being linear-time. However, when α is not constant, this is not the case anymore; in the worst case where $\alpha \simeq m$, the running times for inversion grow like m^3 (neglecting logarithmic factors again), whereas dense linear algebra algorithms take time $O(m^\omega)$. In [4], Bostan, Jeannerod and Schost gave algorithms for inversion of structured matrices of the four classical types with running time $O(\alpha^{\omega-1} \mathcal{M}(m) \lg(m)^2)$. This is satisfactory for $\alpha \simeq m$, since we recover the cost of dense linear algebra, up to logarithmic factors. However, for $\alpha = O(1)$, this algorithm is slightly slower (by a factor $\lg(m)$) than the MBA algorithm.

Main results. The results in this paper cover in a uniform manner the general class of operators based on companion matrices introduced above. Inspired by [4], we obtain algorithms for structured matrix multiplication, inversion, and linear system solving whose running times grow with α as $\alpha^{\omega-1}$ instead of α^2 ; however, we manage to avoid the loss of the $\lg(m)$ factor observed in [4], thus improving on the results of that paper. All our algorithms assume exact arithmetic (and we do not claim any kind of accuracy guarantee when using finite precision arithmetic). Note also that they are deterministic for multiplication and, similarly to [4], randomized (Las Vegas) for inversion and linear system solving.

Let $\mathbf{P}$ and $\mathbf{Q}$ be as before. We give in Section 3 inversion formulas for the basic operators $\mathcal{L} = \nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ or $\mathcal{L} = \Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ that allow us to recover a matrix $\mathbf{A}$ from one of its generators, generalizing simultaneously previous results for Toeplitz, Hankel, Vandermonde, and Cauchy displacement operators. From these formulas, one readily deduces that a matrix-vector product $\mathbf{A}\mathbf{u}$ can be computed in time

$$\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, 1) = O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha \mathcal{E}(\mathbf{P}, \mathbf{Q})),$$

which is $\tilde{O}(\alpha p)$ for $p = \max(m, n)$.

For matrix-matrix products, the direct approach which is based on the estimate $\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) \leq \beta \mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, 1)$ thus leads to $\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) = \tilde{O}(\alpha \beta p)$, which is sub-optimal, as we pointed out before. The following theorem improves on this straightforward approach, for all operators associated with $(\mathbf{P}, \mathbf{Q})$.

THEOREM 1. *Let $\mathbf{P} = P_1, \dots, P_d$ and $\mathbf{Q} = Q_1, \dots, Q_e$ be two families of polynomials in $\mathbb{F}[x]$. Assume that in each of these families the polynomials are monic, non-constant and pairwise coprime, and denote $m = \sum_{i=1}^d \deg(P_i)$ and $n = \sum_{i=1}^e \deg(Q_i)$. Then, for any invertible operator $\mathcal{L}$ associated with $(\mathbf{P}, \mathbf{Q})$, we can take*

$$\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) = O\left(\frac{\beta'}{\alpha'} \mathcal{M}'_{\text{mat}}\left(\frac{p}{\alpha'}, \alpha'\right) + \mathcal{D}(\mathbf{P}, \mathbf{Q}) + \beta' \mathcal{E}(\mathbf{P}, \mathbf{Q})\right),$$

with $p = \max(m, n)$, $\alpha' = \min(\alpha, \beta)$, and $\beta' = \max(\alpha, \beta)$. Furthermore,

$$\mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha), \mathcal{C}_{\text{solve}}(\mathcal{L}, \alpha) = O\left(\mathcal{M}_{\text{mat}}''\left(\frac{p}{\alpha}, \alpha\right)\right).$$

Using the estimates for $\mathcal{M}_{\text{mat}}'$ given in (4), the estimate for $\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta)$ becomes

$$\begin{cases} O\left(\alpha'^{\omega-2}\beta'\mathcal{M}(p) + \beta'\mathcal{M}(p)\lg(p)\right) & \text{in general,} \\ O\left(\alpha'^{\omega-2}\beta'p + \beta'\mathcal{M}(p)\lg(p)\right) & \text{if } \text{char}(\mathbb{F}) = 0 \text{ or } |\mathbb{F}| \geq 2p. \end{cases}$$

Thus, these results provide a continuum between two extreme cases. When $\alpha = \beta = O(1)$, the cost is $O(\mathcal{M}(p)\lg(p))$; when α and β are large, the first term is dominant, with total costs respectively $O(\alpha'^{\omega-2}\beta'\mathcal{M}(p))$ and $O(\alpha'^{\omega-2}\beta'p)$. Disregarding logarithmic factors, this is always $\tilde{O}(\alpha'^{\omega-2}\beta'p)$: this matches the cost (up to logarithmic factors) of dense, unstructured linear algebra algorithms for multiplying matrices of dimensions $\alpha' \times p$ and $p \times \beta'$ with $\alpha' \leq \min(p, \beta')$.

In fact, when $\alpha = p \leq \beta$ we recover exactly the cost bound $O(\beta p^{\omega-1})$ of the multiplication of two dense unstructured matrices of dimensions $p \times p$ and $p \times \beta$, since then $\mathcal{C}_{\text{mul}}(\mathcal{L}, p, \beta) = O(\frac{\beta}{p}\mathcal{M}_{\text{mat}}'(1, p) + \mathbf{M}(p)\lg(p))$ and $\mathcal{M}_{\text{mat}}'(1, p) = O(p^\omega)$.

For inversion and system solving, the bounds given above on $\mathcal{M}_{\text{mat}}''$ yield

$$\begin{cases} O(\alpha^{\omega-1}\mathcal{M}(p)\lg(p)) & \text{in general,} \\ O(\alpha^{\omega-1}p\lg(p) + \alpha\lg(\alpha)\mathcal{M}(p)\lg(p)) & \text{if } \text{char}(\mathbb{F}) = 0 \text{ or } |\mathbb{F}| \geq 2p; \end{cases}$$

and, when $\alpha = p$ (so that the input matrix is $p \times p$ and not structured with respect to the operator $\mathcal{L}$), the obtained cost is $O(\mathcal{M}_{\text{mat}}''(1, p)) \subset O(p^\omega)$, as can be expected.

The following theorem highlights situations where we can take both $\mathcal{C}(\mathbf{P}, \mathbf{Q})$ and $\mathcal{D}(\mathbf{P}, \mathbf{Q})$ in $O(\mathcal{M}(p))$. In this case, we get slightly better bounds than in the general case for the problems `mul` (for inversion and solve, the terms above are always negligible, for any choice of $\mathbf{P}$ and $\mathbf{Q}$). In view of Tables 1 and 2, we see that these special cases correspond to operators of Toeplitz / Hankel type, or Vandermonde / Cauchy type, when their supports are points in geometric progression.

THEOREM 2. *All notation and assumptions being as in Theorem 1, suppose in addition that one of the following holds:*

- either $d = 1$ and $P_1 = x^m - \varphi$ for some $\varphi \in \mathbb{F}$, or $d = m$ and there exist $u, q \in \mathbb{F}$ such that $P_i = x - uq^i$ for all i ;
- either $e = 1$ and $Q_1 = x^n - \psi$ for some $\psi \in \mathbb{F}$, or $e = n$ and there exist $v, r \in \mathbb{F}$ such that $Q_j = x - vr^j$ for all j .

Then, for any invertible operator $\mathcal{L}$ associated with $(\mathbf{P}, \mathbf{Q})$, we can take

$$\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) = O\left(\frac{\beta'}{\alpha'}\mathcal{M}_{\text{mat}}'\left(\frac{p}{\alpha'}, \alpha'\right) + \beta'\mathcal{M}(p)\right),$$

with $p = \max(m, n)$, $\alpha' = \min(\alpha, \beta)$, and $\beta' = \max(\alpha, \beta)$.

Using once again the bounds on $\mathcal{M}_{\text{mat}}'$ given in (4), this is thus

$$\begin{cases} O(\alpha'^{\omega-2}\beta'\mathcal{M}(p)) & \text{in general,} \\ O(\alpha'^{\omega-2}\beta'p + \beta'\lg(\alpha')\mathcal{M}(p)) & \text{if } \text{char}(\mathbb{F}) = 0 \text{ or } |\mathbb{F}| \geq 2p. \end{cases}$$

Comparison with previous work. Let us briefly compare our results with previous ones, in increasing order of generality.

For classical operators of Hankel, Toeplitz, Vandermonde or Cauchy type, our results match classical ones in the cases $\alpha = \beta = O(1)$ (for multiplication) or $\alpha = O(1)$ (for inversion). When we drop such assumptions, our results improve on all previous ones. For instance, for the inversion of Toeplitz-like or Hankel-like matrices, the best previous results were either $O(\alpha^2 \mathcal{M}(m) \lg(m))$ in [25, 3, 19], or $O(\alpha^{\omega-1} \mathcal{M}(m) \lg(m)^2)$ in [4]. We improve them simultaneously, by dropping the cost to $O(\alpha^{\omega-1} \mathcal{M}(m) \lg(m))$.

We mentioned earlier the work of Olshevsky and Shokrollahi [27, 28], who consider Sylvester operators where the displacement matrices are block-diagonal, with Jordan blocks. A Jordan block of size m associated with $\lambda \in \mathbb{F}$ is similar to the companion matrix of $(x - \lambda)^m$; the similarity matrix is the “Pascal” matrix of the mapping $F(x) \mapsto F(x + \lambda)$, for $F \in \mathbb{F}[x]_m$. Because the similarity matrix and its inverse can be applied in time $O(\mathcal{M}(m))$ to a vector [1], the problem considered by Olshevsky and Shokrollahi is essentially equivalent to a particular case of ours.

As it turns out, for these particular cases, the results in [27, 28] for matrix/vector product and inversion are equivalent to ours, when the displacement rank α is $O(1)$. For larger α , our results improve on those of [27, 28], whose costs are quadratic in α .

Finally, to the best of our knowledge, no previous work addressed the general case of block-companion matrices that we deal with in this paper.

An application. We are able to address problems related to *simultaneous approximation* using our operators. Suppose that we are given $\mathbf{P} = P_1, \dots, P_d$, with sum of degrees m , together with “residuals”

$$R_{1,1}, \dots, R_{1,\alpha}, \dots, R_{d,1}, \dots, R_{d,\alpha},$$

with all $R_{i,j}$ in $\mathbb{F}[x]$ such that $\deg(R_{i,j}) < \deg(P_i)$ for all i, j . We are looking for polynomials $f_1, \dots, f_\alpha$ in $\mathbb{F}[x]$ such that $f_1 R_{i,1} + \dots + f_\alpha R_{i,\alpha} = 0 \bmod P_i$ holds for $i = 1, \dots, d$, and with prescribed degree bounds (the polynomial f_j should be in $\mathbb{F}[x]_{n_j}$, with $n_1 + \dots + n_\alpha = O(m)$).

The matrix of the corresponding linear system, where the coefficients of $f_1, \dots, f_\alpha$ are unknown, has size $O(m)$ and displacement rank $O(\alpha)$ for the operator $\Delta_{\mathbf{M}_P, \mathbf{M}_Q^t}$, where $\mathbf{Q}$ is the “vector” consisting of the unique polynomial $Q_1 = x^m - \varphi$, where φ is chosen such that the assumptions of Theorem 1 hold. Hence, we can solve such a system using $O(\alpha^{\omega-1} \mathcal{M}(m) \log(m))$ base field operations. This is to be compared with an algorithm given in [10], that has cost $O((\alpha + d)^{\omega-1} \mathcal{M}(m) \log(m)^2)$, but which does not require the assumption that $P_1, \dots, P_d$ be pairwise coprime.

This problem generalizes Hermite-Padé (with $d = 1$ and $P_1 = x^m$) and so-called M-Padé approximation, with $P_i = (x - x_i)^{m_i}$, for pairwise distinct x_i in $\mathbb{F}$. A typical instance is the reconstruction of the minimal polynomial of an algebraic function. Suppose that f is a root of a polynomial $S \in \mathbb{F}[x][y]$, so f lies in the algebraic closure of $\mathbb{F}(x)$. Given the power series expansion of f at high enough precision around a point x_1 , it is possible to reconstruct S by means of Hermite-Padé approximation. Using our algorithm, we obtain the same output for the same cost, starting this time from approximations at points $x_1, \dots, x_d$, each of them requiring smaller precision.

Organization of the paper. Section 2 introduces notation used all along this paper, and discusses a few classical questions about polynomial arithmetic, such as Chinese remaindering. Section 3 gives *inversion formulas* for the operators we discuss, which generalize well-known results for classical operators. In Section 4, we use these formulas in order to reduce most of our questions to similar questions for *basic operators*, or

even Toeplitz / Hankel operators. Section 5 gives the main algorithmic result in this paper, an improved algorithm for the multiplication of structured matrices. Finally, in Section 6, we show how this new multiplication algorithm can be used within the MBA approach to accelerate structured inversion and structured system solving.

Acknowledgements. We thank an anonymous referee for many detailed comments.

2. Preliminaries. In this section, we review some classical operations on polynomials and matrices: we introduce a short list of useful matrices, such as multiplication matrices, reversal matrices and Krylov matrices, and we discuss questions related to modular computations, or multiple reduction and Chinese remaindering.

Throughout the article, we will use some well-known complexity results on polynomial arithmetic; as a general rule, they can be found in the books [12], [8], or [33].

2.1. Basic notation. Matrices (resp. vectors) are written in upper-case (resp. lower-case) sans-serif font. If A (resp. $B, C, \dots$) is a matrix, $\mathbf{a}_i$ (resp. $\mathbf{b}_i, \mathbf{c}_i, \dots$) is its i th column. If $\mathbf{x}$ (resp. $\mathbf{y}, \mathbf{z}, \dots$) is a vector, its i th entry is written x_i (resp. $y_i, z_i, \dots$). Special matrices (cyclic, companion, diagonal, ...) will be written with blackboard bold letters ($\mathbb{Z}, \mathbb{M}, \mathbb{D}, \dots$). The entries of an $m \times n$ matrix A are indexed from 0 to $m - 1$ (row indices) and from 0 to $n - 1$ (column indices).

We will also use the following notation.

- For $\mathbf{u} = [u_0 \ \dots \ u_{m-1}]^t \in \mathbb{F}^m$, we write $\text{pol}(\mathbf{u})$ to denote the polynomial $u_0 + \dots + u_{m-1}x^{m-1} \in \mathbb{F}[x]_m$.
- For any polynomial $F \in \mathbb{F}[x]$ of degree at most d , $\text{rev}(F, d)$ denotes its reverse polynomial $x^d F(1/x) \in \mathbb{F}[x]$.

For $m \geq 1$, $\mathbb{J}_m$ is the $m \times m$ reversal matrix, with 1s on the antidiagonal only. More generally, for $1 \leq \ell \leq m$, $\mathbb{J}_{\ell, m}$ is the $m \times m$ matrix with 1s on the upper antidiagonal entries of indices $(\ell - 1, 0), \dots, (0, \ell - 1)$, so $\mathbb{J}_{m, m} = \mathbb{J}_m$; by convention, for $\ell = 0$, $\mathbb{J}_{0, m}$ is the zero matrix of size m .

Finally, for $A \in \mathbb{F}^{m \times m}$, $\mathbf{v} \in \mathbb{F}^m$, and $\ell \geq 1$, we write $\mathbb{K}(A, \mathbf{v}, \ell)$ to denote the Krylov matrix in $\mathbb{F}^{m \times \ell}$ whose columns are $\mathbf{v}, A\mathbf{v}, \dots, A^{\ell-1}\mathbf{v}$.

2.2. Chinese remaindering and related problems. In this subsection, we introduce additional notation and give basic results related to multiple reduction and Chinese remaindering. Consider pairwise-coprime monic polynomials $\mathbf{P} = P_1, \dots, P_d$ in $\mathbb{F}[x]$ with $\deg(P_i) = m_i$, and let $m = m_1 + \dots + m_d$ and $P = P_1 \dots P_d$.

Recall that we associated with the family $\mathbf{P}$ the cost functions $\mathcal{C}(\mathbf{P})$ and $\mathcal{D}(\mathbf{P})$. They will help us measure the cost of the following operations. To $\mathbf{P}$, we associate the *multiple reduction* mapping, and its inverse, *Chinese remaindering*, defined by

$$\begin{aligned} \text{red}_{\mathbf{P}} : \quad \mathbb{F}[x]_m &\rightarrow \mathbb{F}[x]_{m_1} \times \dots \times \mathbb{F}[x]_{m_d} \\ A &\mapsto (A \bmod P_1, \dots, A \bmod P_d) \end{aligned}$$

and

$$\begin{aligned} \text{crt}_{\mathbf{P}} : \quad \mathbb{F}[x]_{m_1} \times \dots \times \mathbb{F}[x]_{m_d} &\rightarrow \mathbb{F}[x]_m \\ (A_1, \dots, A_d) &\mapsto A, \text{ such that } A \bmod P_i = A_i \text{ for all } i. \end{aligned}$$

A related question is *linear recombination*, defined as the following isomorphism:

$$\begin{aligned} \text{comb}_{\mathbf{P}} : \quad \mathbb{F}[x]_{m_1} \times \dots \times \mathbb{F}[x]_{m_d} &\rightarrow \mathbb{F}[x]_m \\ (A_1, \dots, A_d) &\mapsto A_1 P_2 \dots P_d + \dots + P_1 \dots P_{d-1} A_d. \end{aligned}$$

Fast algorithms for these three operations lead to the costs given in the next lemma. Although such costs can be found in or deduced from [8, 12], they do not seem to have

been presented in this way, with a common precomputation involving both functions $\mathcal{C}(\mathbf{P})$ and $\mathcal{D}(\mathbf{P})$, and then an extra cost involving only $\mathcal{C}(\mathbf{P})$.

LEMMA 3. *Let $\mathbf{P} = P_1, \dots, P_d$ as above be given and assume that $\mathcal{H}_{\mathbf{P}}$ holds. Then, after a precomputation of time $O(\mathcal{C}(\mathbf{P}) + \mathcal{D}(\mathbf{P}))$ that yields P as a by-product, one can apply the maps $\text{red}_{\mathbf{P}}$, its inverse $\text{crt}_{\mathbf{P}}$, as well as $\text{comb}_{\mathbf{P}}$ and its inverse, to any vector in time $O(\mathcal{C}(\mathbf{P}))$.*

Proof. For $1 \leq i \leq d$, let us define $E_i = P/P_i \bmod P_i$ and $F_i = 1/E_i \bmod P_i$. (In particular, if $d = 1$ then $E_1 = F_1 = 1$, while if P_i is the linear polynomial $x - x_i$ then $E_i = 1/F_i$ is the value of the derivative of P at x_i .) Using these polynomials, we first note how the mappings $\text{crt}_{\mathbf{P}}$ and $\text{comb}_{\mathbf{P}}$ relate to each other: Chinese remaindering of $A_1, \dots, A_d$ is done by computing $A_i F_i \bmod P_i$ for all i , and then applying $\text{comb}_{\mathbf{P}}$ to these products; conversely, to perform linear recombination it suffices to compute the products $A_i E_i \bmod P_i$ and then to apply $\text{crt}_{\mathbf{P}}$. Note also that the inverse of $\text{comb}_{\mathbf{P}}$ can be obtained by first using $\text{red}_{\mathbf{P}}$ and then multiplying by the F_i 's modulo P_i .

Suppose we are in the case where $P_i = x - x_i$ with the x_i 's in *geometric* progression. Then, we can compute $P, E_1, \dots, E_d$ and apply $\text{red}_{\mathbf{P}}$ and $\text{crt}_{\mathbf{P}}$ in time $O(\mathcal{C}(\mathbf{P}))$ using the algorithms of [7]. Using the remark above on the relation between $\text{crt}_{\mathbf{P}}$ and $\text{comb}_{\mathbf{P}}$, the claim carries over to $\text{comb}_{\mathbf{P}}$ and its inverse, so the lemma is proved in this case.

In the general case, we can precompute P , as well as apply $\text{red}_{\mathbf{P}}$ and $\text{comb}_{\mathbf{P}}$ (without any further precomputation) in time $O(\mathcal{C}(\mathbf{P}))$ using respectively Theorems 2.19, 3.19 and Step 2 (or 5) of Theorem 3.21 in [8].

Let $P^* = P/P_1 + \dots + P/P_d$. Since P^* is obtained by applying $\text{comb}_{\mathbf{P}}$ to the polynomials $(1, \dots, 1)$, it can thus be precomputed in time $O(\mathcal{C}(\mathbf{P}))$. The modular images $P^* \bmod P_i$, which coincide with the polynomials E_i seen above, can be precomputed in the same time by applying $\text{red}_{\mathbf{P}}$. Finally, we precompute the inverses $F_i = 1/E_i \bmod P_i$ in time $O(\mathcal{D}(\mathbf{P}))$ by fast extended GCD (when $d = 1$, this is straightforward, since we have seen that $E_1 = 1$; otherwise, each F_i is obtained in time $O(\mathcal{M}(m_i) \lg(m_i))$ using the fast extended Euclidean algorithm [8, Cor. 3.14]).

Using the first paragraph, we can then perform Chinese remaindering for the cost $O(\mathcal{C}(\mathbf{P}))$ of $\text{comb}_{\mathbf{P}}$, plus modular multiplications by the F_i 's. The cost of the latter is $O(\mathcal{M}(m))$, by the super-linearity of $\mathcal{M}$, which is $O(\mathcal{C}(\mathbf{P}))$. The same holds for the inverse of $\text{comb}_{\mathbf{P}}$, which is reduced to $\text{red}_{\mathbf{P}}$ and modular multiplications by the F_i 's. $\square$

In matrix terms, we will need the following notation in the next sections (here and hereafter, we use canonical monomial bases for vector spaces such as $\mathbb{F}[x]_m$):

- $\mathbb{X}_{\mathbf{P}} \in \mathbb{F}^{m \times m}$ is the matrix of the mapping $\text{comb}_{\mathbf{P}}$;
- $\mathbb{W}_{\mathbf{P}} \in \mathbb{F}^{m \times m}$ is the matrix of the mapping $\text{red}_{\mathbf{P}}$.

LEMMA 4. *Let $\mathbf{P} = P_1, \dots, P_d$ as above be given and assume that $\mathcal{H}_{\mathbf{P}}$ holds. Then, after a precomputation of time $O(\mathcal{C}(\mathbf{P}) + \mathcal{D}(\mathbf{P}))$, one can compute $\mathbb{X}_{\mathbf{P}} \mathbf{u}$, as well as $\mathbb{W}_{\mathbf{P}} \mathbf{u}$, $\mathbb{W}_{\mathbf{P}}^t \mathbf{u}$, $\mathbb{W}_{\mathbf{P}}^{-1} \mathbf{u}$ and $\mathbb{W}_{\mathbf{P}}^{-t} \mathbf{u}$ for any $\mathbf{u}$ in $\mathbb{F}^m$ in time $O(\mathcal{C}(\mathbf{P}))$.*

Proof. For $\mathbb{X}_{\mathbf{P}} \mathbf{u}$, $\mathbb{W}_{\mathbf{P}} \mathbf{u}$ and $\mathbb{W}_{\mathbf{P}}^{-1} \mathbf{u}$, this is nothing else than the previous lemma. To compute with the transpose of $\mathbb{W}_{\mathbf{P}}$, we could rely on the transposition principle [6], but it is actually possible to give direct algorithms. Algorithm `TSimulMod` in [5] shows how to reduce the computation of $\mathbb{W}_{\mathbf{P}}^t \mathbf{u}$ to an application of $\text{comb}_{\mathbf{P}}$ and a few power series multiplications / inversions that involve only P and the P_i 's, and whose costs add up to $O(\mathcal{M}(m))$. Thus, the claimed $O(\mathcal{C}(\mathbf{P}))$ follows from the previous lemma. It is straightforward to invert that algorithm step-by-step, obtaining an algorithm for computing $\mathbb{W}_{\mathbf{P}}^{-t} \mathbf{u}$ in time $O(\mathcal{C}(\mathbf{P}))$ as well. $\square$

2.3. Modular arithmetic. Given a monic polynomial $P \in \mathbb{F}[x]$ of degree m , we will frequently work with the residue class ring $\mathbb{A} = \mathbb{F}[x]/P$; doing so, we will identify its elements with polynomials in $\mathbb{F}[x]_m$. In this subsection, we review several questions related to computations in $\mathbb{A}$, such as modular multiplication and its transpose.

As a preamble, we define two useful matrices related to P . First, for $\ell \geq 1$, $\mathbb{W}_{P,\ell} \in \mathbb{F}^{m \times \ell}$ denotes the matrix of reduction modulo P , which maps a polynomial $F \in \mathbb{F}[x]_\ell$ to $F \bmod P$. Computationally, applying $\mathbb{W}_{P,\ell}$ to a vector amounts to doing a Euclidean division. Next, writing $P = p_0 + p_1x + \cdots + p_mx^m$ (so that $p_m = 1$), we will denote by $\mathbb{Y}_P$ the $m \times m$ Hankel matrix

$$\mathbb{Y}_P = \begin{bmatrix} p_1 & \cdots & p_{m-1} & 1 \\ \vdots & \ddots & 1 & \\ p_{m-1} & \ddots & & \\ 1 & & & \end{bmatrix}.$$

Operations with $\mathbb{Y}_P$ are fast, due to the triangular Hankel nature of this matrix. The only non-trivial part of the following lemma concerns the inverse of $\mathbb{Y}_P$. It is proved in [33, §2.5] for triangular Toeplitz matrices; the extension to the Hankel case is straightforward, since $\mathbb{Y}_P^{-1} = (\mathbb{J}_m \mathbb{Y}_P)^{-1} \mathbb{J}_m$ with $\mathbb{J}_m \mathbb{Y}_P$ triangular Toeplitz.

LEMMA 5. *Given P , for $\mathbf{u}$ in $\mathbb{F}^m$, one can compute $\mathbb{Y}_P \mathbf{u}$ and $\mathbb{Y}_P^{-1} \mathbf{u}$ in $O(\mathcal{M}(m))$.*

Let us now discuss multiplication in $\mathbb{A}$, often called *modular multiplication*. Computationally, it is done by a standard polynomial multiplication, followed by a Euclidean division; altogether, this takes time $O(\mathcal{M}(m))$.

In matrix terms, recall that $\mathbb{M}_P$ denotes the $m \times m$ companion matrix associated with P ; equivalently, this is the matrix of the multiplication-by- x endomorphism in $\mathbb{A}$. More generally, for F in $\mathbb{F}[x]$, we denote by $\mathbb{M}_{F,P}$ the $m \times m$ matrix of multiplication by F in $\mathbb{A}$; that is, the matrix whose (i, j) entry is the coefficient of x^i in $Fx^j \bmod P$, for $i, j = 0, \dots, m-1$. Thus, by what was said above, applying $\mathbb{M}_{F,P}$ to a vector can be done in time $O(\mathcal{M}(m))$. Note also that $\mathbb{M}_{x,P} = \mathbb{M}_P$ and that $\mathbb{M}_{F,P} = F(\mathbb{M}_P)$.

Yet more generally, for $\ell \geq 1$, $\mathbb{M}_{F,P,\ell} \in \mathbb{F}^{m \times \ell}$ is the matrix whose entries are the coefficients of the same polynomials $Fx^j \bmod P$ as above, but now for $j = 0, \dots, \ell-1$. The following easy lemma simply says that for F in $\mathbb{F}[x]_m$ and G in $\mathbb{F}[x]_\ell$, $(GF) \bmod P = ((G \bmod P)F) \bmod P$.

LEMMA 6. *Let F be in $\mathbb{F}[x]_m$ and $\ell \geq 1$. Then $\mathbb{M}_{F,P,\ell} = \mathbb{M}_{F,P} \mathbb{W}_{P,\ell}$.*

We will also use a dual operation, involving the space $\mathbb{A}^*$ of $\mathbb{F}$ -linear forms $\mathbb{A} \rightarrow \mathbb{F}$. Such a linear form will be given by the values it takes on the monomial basis $1, x, \dots, x^{m-1}$ of $\mathbb{A}$. Then, *transposed multiplication* is the operation mapping $(F, \lambda) \in \mathbb{A} \times \mathbb{A}^*$ to the linear form $\lambda' = F \circ \lambda$, defined by $\lambda'(x^i) = \lambda(x^i F)$, where the multiplication takes place in $\mathbb{A}$. The name reflects the fact that for fixed F , transposed multiplication by F is indeed the dual map of the multiplication-by- F endomorphism of $\mathbb{A}$. In matrix terms, transposed product by F amounts to multiplication by $\mathbb{M}_{F,P}^t$.

Transposed products can be computed in the same time $O(\mathcal{M}(m))$ as “standard” modular multiplications, by a dual form of modular multiplication [6]. However, such an algorithm relies on middle product techniques [15], which are not straightforward to describe. The following lemma shows an alternative approach, with same cost up to a constant factor: to perform a transposed product by F , it suffices to do a modular multiplication by F , with a pre- and post-multiplication by $\mathbb{Y}_P$ and its inverse (which

can both be done in $O(\mathcal{M}(m))$, see Lemma 5). For $F = x$, this result is sometimes referred to as saying that $\mathbb{Y}_P$ is a *symmetrizer* for the polynomial P ; see [22, p. 455].

LEMMA 7. *For all $F \in \mathbb{F}[x]_m$, we have $\mathbb{Y}_P \mathbb{M}_{F,P}^t = \mathbb{M}_{F,P} \mathbb{Y}_P$. Equivalently, $\mathbb{M}_{F,P}^t = \mathbb{Y}_P^{-1} \mathbb{M}_{F,P} \mathbb{Y}_P$.*

Proof. By linearity, it is enough to consider the case of $F = x^\ell$, for $0 \leq \ell < m$. The case $\ell = 0$ is clear, since $\mathbb{M}_{1,P}$ is the identity matrix. For $\ell = 1$, this result is well-known, see for instance [22, Ex. 3, Ch. 13]; it can be proved by simple inspection. Once the claim is established for x , an easy induction shows that it holds for x^ℓ , for an arbitrary $\ell \geq 1$, using the fact that $\mathbb{M}_{x^{\ell+1},P} = \mathbb{M}_P \mathbb{M}_{x^\ell,P} = \mathbb{M}_{x^\ell,P} \mathbb{M}_P$. $\square$

We now discuss Krylov matrices derived from (transposed) companion matrices.

LEMMA 8. *Let $\mathbf{v} \in \mathbb{F}^m$, $\ell \geq 1$, and $F = \text{pol}(\mathbf{v}) \in \mathbb{F}[x]_m$. Then*

- $\mathbb{K}(\mathbb{M}_P, \mathbf{v}, \ell) = \mathbb{M}_{F,P} \mathbb{W}_{P,\ell} = \mathbb{M}_{F,P,\ell}$ and
- $\mathbb{K}(\mathbb{M}_P^t, \mathbf{v}, \ell) = \mathbb{Y}_P^{-1} \mathbb{M}_{G,P} \mathbb{W}_{P,\ell}$ with $G = \text{pol}(\mathbb{Y}_P \mathbf{v})$.

Proof. The first assertion is clear. Indeed, the columns of the Krylov matrix $\mathbb{K}(\mathbb{M}_P, \mathbf{v}, \ell)$ are the coefficient vectors of the polynomials $x^j F \bmod P$, for $0 \leq j < \ell$, so that $\mathbb{K}(\mathbb{M}_P, \mathbf{v}, \ell) = \mathbb{M}_{F,P,\ell}$; the claim now follows from Lemma 6. For the second assertion, the fact that $\mathbb{M}_P^t = \mathbb{Y}_P^{-1} \mathbb{M}_P \mathbb{Y}_P$, which follows from Lemma 7, implies that $\mathbb{K}(\mathbb{M}_P^t, \mathbf{v}, \ell) = \mathbb{Y}_P^{-1} \mathbb{K}(\mathbb{M}_P, \mathbb{Y}_P \mathbf{v}, \ell)$. Using the first point concludes the proof. $\square$

2.4. Computations with a family of polynomials. We now consider a family of monic polynomials $\mathbf{P} = P_1, \dots, P_d$, and briefly discuss the extension of the previous claims to this context. As before, we write $m_i = \deg(P_i)$ and $m = m_1 + \dots + m_d$.

In terms of matrices, we have already mentioned the definition of the block-diagonal companion matrix $\mathbb{M}_{\mathbf{P}}$ associated with $\mathbf{P}$, whose blocks are $\mathbb{M}_{P_1}, \dots, \mathbb{M}_{P_d}$. Similarly, $\mathbb{Y}_{\mathbf{P}}$ will denote the block-diagonal matrix with Hankel blocks $\mathbb{Y}_{P_1}, \dots, \mathbb{Y}_{P_d}$. The next lemma, which is about $\mathbb{Y}_{\mathbf{P}}$ and its inverse, is a direct consequence of Lemma 5; the complexity estimate follows from the super-linearity of the function $\mathcal{M}$.

LEMMA 9. *Given $\mathbf{P}$, for $\mathbf{u}$ in $\mathbb{F}^m$, one can compute $\mathbb{Y}_{\mathbf{P}} \mathbf{u}$ and $\mathbb{Y}_{\mathbf{P}}^{-1} \mathbf{u}$ in $O(\mathcal{M}(m))$.*

Another straightforward result is the following extension of Lemma 7.

LEMMA 10. *The relation $\mathbb{Y}_{\mathbf{P}} \mathbb{M}_{\mathbf{P}}^t = \mathbb{M}_{\mathbf{P}} \mathbb{Y}_{\mathbf{P}}$ holds.*

Next, we discuss computations with Krylov matrices derived from $\mathbb{M}_{\mathbf{P}}$ and its transpose. We will only need some special cases, for which very simple formulas are available (and for which invertibility will be easy to prove). Recall that $\mathbb{W}_{\mathbf{P}} \in \mathbb{F}^{m \times m}$ denotes the matrix of the map $\text{red}_{\mathbf{P}}$ of multiple reduction modulo $P_1, \dots, P_d$; thus, it is obtained by stacking the matrices $\mathbb{W}_{P_1,m}, \dots, \mathbb{W}_{P_d,m}$.

LEMMA 11. *Writing $\mathbf{e}_{m,i}$ to denote the i th unit vector in $\mathbb{F}^m$, the following holds:*

- $\mathbb{K}(\mathbb{M}_{\mathbf{P}}, \mathbf{v}, m) = \mathbb{W}_{\mathbf{P}}$ with $\mathbf{v} = [\mathbf{e}_{m_1,1}^t | \dots | \mathbf{e}_{m_d,1}^t]^t \in \mathbb{F}^m$, and
- $\mathbb{K}(\mathbb{M}_{\mathbf{P}}^t, \mathbf{w}, m) = \mathbb{Y}_{\mathbf{P}}^{-1} \mathbb{W}_{\mathbf{P}}$ with $\mathbf{w} = [\mathbf{e}_{m_1,m_1}^t | \dots | \mathbf{e}_{m_d,m_d}^t]^t \in \mathbb{F}^m$.

Proof. The block structure of $\mathbb{M}_{\mathbf{P}}$ and the definition of $\mathbf{v}$ imply that $\mathbb{K}(\mathbb{M}_{\mathbf{P}}, \mathbf{v}, m)$ is obtained by stacking the matrices $\mathbb{K}(\mathbb{M}_{P_i}, \mathbf{e}_{m_i,1}, m)$ for $i = 1, \dots, d$. Since $\text{pol}(\mathbf{e}_{m_i,1}) = 1 \in \mathbb{F}[x]_{m_i}$, Lemma 8 gives $\mathbb{K}(\mathbb{M}_{P_i}, \mathbf{e}_{m_i,1}, m) = \mathbb{W}_{P_i,m}$ for all i ; this proves the first case. In the second case, $\mathbb{K}(\mathbb{M}_{\mathbf{P}}^t, \mathbf{w}, m)$ is decomposed into blocks $\mathbb{K}(\mathbb{M}_{P_i}^t, \mathbf{e}_{m_i,m_i}, m)$, which, by Lemma 8 and since $\text{pol}(\mathbb{Y}_{P_i} \mathbf{e}_{m_i,m_i}) = 1 \in \mathbb{F}[x]_{m_i}$, are equal to $\mathbb{Y}_{P_i}^{-1} \mathbb{W}_{P_i,m}$. $\square$

Finally, the following lemma discusses computations involving two families of monic polynomials $\mathbf{P}$ and $\mathbf{Q}$.

LEMMA 12. Let $\mathbf{P} = P_1, \dots, P_d$ and $\mathbf{Q} = Q_1, \dots, Q_e$, such that $\mathcal{H}_{\mathbf{P}}$ and $\mathcal{H}_{\mathbf{Q}}$ hold, and let $P = P_1 \cdots P_d$ and $Q = Q_1 \cdots Q_e$. Then, the following holds:

- If $\gcd(P, Q) = 1$, we can compute all $1/Q \bmod P_i$, for $i = 1, \dots, d$, in time $O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \mathcal{C}(\mathbf{P}, \mathbf{Q}))$.
- Let $n = \deg(Q)$ and $\tilde{Q} = \text{rev}(Q, n)$. If $\gcd(P, \tilde{Q}) = 1$, we can compute all $1/\tilde{Q} \bmod P_i$, for $i = 1, \dots, d$, in time $O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \mathcal{C}(\mathbf{P}, \mathbf{Q}))$.

Proof. Let $p = \max(m, n)$. We first discuss a particular case, where $d = e = 1$ and $P = x^m - \varphi$ and $Q = x^n - \psi$. In this case, we can compute $1/Q \bmod P$ in linear time $O(p)$ [39, Lemma 3], provided it is well-defined; the same holds for $1/\tilde{Q} \bmod P$. Thus, our claim holds, since in this case, $\mathcal{D}(\mathbf{P}, \mathbf{Q}) = m + n$.

If P is not as above, we compute Q in time $O(\mathcal{C}(\mathbf{Q}) + \mathcal{D}(\mathbf{Q}))$ using Lemma 3, then $Q \bmod P$ in time $O(\mathcal{M}(p))$, which is $O(\mathcal{C}(\mathbf{P}, \mathbf{Q}))$. Using Lemma 3 again, we deduce all $Q \bmod P_i$, for $i \leq d$, in time $O(\mathcal{C}(\mathbf{P}))$; finally, we invert all remainders using fast extended GCD with the P_i 's in time $O(\sum_{i \leq d} \mathcal{M}(m_i) \lg(m_i))$; under our assumption on P , this is $O(\mathcal{D}(\mathbf{P}))$. The total cost is $O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \mathcal{C}(\mathbf{P}, \mathbf{Q}))$, so our claim for the inverses of Q modulo the P_i 's holds. We proceed similarly for $\tilde{Q}$.

The last case to consider is when $d = 1$ and $P = x^m - \varphi$, but with Q not of the form $x^n - \psi$. We proceed in the converse direction: we first reduce and invert P modulo all Q_j , for all $j \leq e$; this takes time $O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \mathcal{C}(\mathbf{P}, \mathbf{Q}))$. By Chinese Remaindering, we obtain $R = 1/P \bmod Q$, without increasing the cost. Knowing R , we deduce $S = 1/Q \bmod P$, since R and S satisfy $RP + SQ = 1$; this costs an extra $O(\mathcal{M}(p))$, which is $O(\mathcal{C}(\mathbf{P}, \mathbf{Q}))$, so our claim is proved. We proceed similarly for $\tilde{Q}$. $\square$

3. Inverting displacement operators. Let us consider $\mathbf{P} = P_1, \dots, P_d$ and $\mathbf{Q} = Q_1, \dots, Q_e$ tuples of monic polynomials in $\mathbb{F}[x]$, and let $m_1, \dots, m_d, n_1, \dots, n_e, m, n, P$ and Q be as before. We assume the coprimality conditions $\mathcal{H}_{\mathbf{P}}$ and $\mathcal{H}_{\mathbf{Q}}$. In this section, we establish inversion formulas for the two basic operators of Sylvester and Stein types associated with $(\mathbf{P}, \mathbf{Q})$. In other words, given two matrices $(\mathbf{G}, \mathbf{H})$ in $\mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha}$, we show how to reconstruct the matrices $\mathbf{A}$ and $\mathbf{A}'$ in $\mathbb{F}^{m \times n}$ such that

$$\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}}^t(\mathbf{A}) = \mathbf{G}\mathbf{H}^t \quad \text{and} \quad \Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}}^t(\mathbf{A}') = \mathbf{G}\mathbf{H}^t,$$

provided the corresponding operators are invertible. We will use the following objects.

- We denote by $\mathbf{g}_1, \dots, \mathbf{g}_{\alpha}$ and $\mathbf{h}_1, \dots, \mathbf{h}_{\alpha}$ the columns of $\mathbf{G}$ and $\mathbf{H}$. Corresponding to the partition of $\mathbb{M}_{\mathbf{P}}$ into blocks, we partition each $\mathbf{g}_k$, $k \leq \alpha$, into smaller column vectors $\mathbf{g}_{1,k}, \dots, \mathbf{g}_{d,k}$; similarly, we partition each $\mathbf{h}_k$ into $\mathbf{h}_{1,k}, \dots, \mathbf{h}_{e,k}$ according to the block structure of $\mathbb{M}_{\mathbf{Q}}^t$. The matrices $\mathbf{G}$ and $\mathbf{H}$ themselves are partitioned into matrices $\mathbf{G}_1, \dots, \mathbf{G}_d$ and $\mathbf{H}_1, \dots, \mathbf{H}_e$; $\mathbf{G}_i$ has size $m_i \times \alpha$ and columns $\mathbf{g}_{i,1}, \dots, \mathbf{g}_{i,\alpha}$, whereas $\mathbf{H}_j$ has size $n_j \times \alpha$ and columns $\mathbf{h}_{j,1}, \dots, \mathbf{h}_{j,\alpha}$.
- For $i = 1, \dots, d$ and $k = 1, \dots, \alpha$, we let $g_{i,k} = \text{pol}(\mathbf{g}_{i,k})$, so that $g_{i,k}$ is in $\mathbb{F}[x]_{m_i}$. Similarly, for $j = 1, \dots, e$, we define $h_{j,k} = \text{pol}(\mathbf{h}_{j,k}) \in \mathbb{F}[x]_{n_j}$. We further let γ_k be the unique polynomial in $\mathbb{F}[x]_m$ such that $\gamma_k \bmod P_i = g_{i,k}$ holds for $i = 1, \dots, d$; in other words, we have

$$\gamma_k = \text{crt}_{\mathbf{P}}(g_{1,k}, \dots, g_{d,k}).$$

The polynomial η_k is defined similarly, replacing $g_{i,k}$ and P_i by $h_{j,k}$ and Q_j .

In the following theorem, recall that the necessary and sufficient condition for the Sylvester operator to be invertible is that $\gcd(P, Q) = 1$; for the Stein operator, the condition is that $\gcd(P, \tilde{Q}) = 1$, with $\tilde{Q} = \text{rev}(Q, n)$.

THEOREM 13. *The following holds:*

- Suppose that $\gcd(P, Q) = 1$. Then, the unique matrix $A \in \mathbb{F}^{m \times n}$ such that $\nabla_{\mathbb{M}_P, \mathbb{M}_Q^t}(A) = GH^t$ is given by

$$A = \mathbb{V}_{P, Q} \mathbb{W}_P \left(\sum_{k \leq \alpha} \mathbb{M}_{\gamma_k, P, n} \mathbb{M}_{\eta_k, Q} \right) \mathbb{X}_Q \mathbb{Y}_Q,$$

where $\mathbb{V}_{P, Q} \in \mathbb{F}^{m \times m}$ is the block-diagonal matrix with blocks $\mathbb{M}_{Q^{-1}, P_i}$ for $i = 1, \dots, d$, where $\mathbb{M}_{Q^{-1}, P_i}$ denotes $\mathbb{M}_{Q^{-1} \bmod P_i, P_i}$.

- Suppose that $\gcd(P, \tilde{Q}) = 1$. Then, the unique matrix $A' \in \mathbb{F}^{m \times n}$ such that $\Delta_{\mathbb{M}_P, \mathbb{M}_Q^t}(A') = GH^t$ is given by

$$A' = \mathbb{V}'_{P, Q} \mathbb{W}_P \left(\sum_{k \leq \alpha} \mathbb{M}_{\gamma_k, P, n} \mathbb{J}_n \mathbb{M}_{\eta_k, Q} \right) \mathbb{X}_Q \mathbb{Y}_Q,$$

where $\mathbb{V}'_{P, Q} \in \mathbb{F}^{m \times m}$ is the block-diagonal matrix with blocks $\mathbb{M}_{\tilde{Q}^{-1}, P_i}$ for $i = 1, \dots, d$, where $\mathbb{M}_{\tilde{Q}^{-1}, P_i}$ denotes $\mathbb{M}_{\tilde{Q}^{-1} \bmod P_i, P_i}$.

The following corollary on the cost of matrix-vector multiplication follows directly; the more difficult case of matrix-matrix multiplication will be handled in Section 5.

COROLLARY 14. *We can take*

$$\mathcal{C}_{\text{mul}}(\nabla_{\mathbb{M}_P, \mathbb{M}_Q^t}, \alpha, 1) = O(\mathcal{D}(P, Q) + \alpha \mathcal{C}(P, Q))$$

and

$$\mathcal{C}_{\text{mul}}(\Delta_{\mathbb{M}_P, \mathbb{M}_Q^t}, \alpha, 1) = O(\mathcal{D}(P, Q) + \alpha \mathcal{C}(P, Q)).$$

Proof. The proof is mostly the same in both cases; it amounts to estimating first the cost of computing the polynomials that define the matrices involved in Theorem 13, and then the cost of multiplying each of these matrices by a single vector.

Given the families of polynomials P and Q , we can compute the products P and Q in time $O(\mathcal{C}(P, Q) + \mathcal{D}(P, Q))$, by Lemma 3. Moreover, since both assumptions $\mathcal{H}_P$ and $\mathcal{H}_Q$ hold and depending on whether $\gcd(P, Q) = 1$ or $\gcd(P, \tilde{Q}) = 1$, we deduce from Lemma 12 that we can compute the inverses of Q , or of $\tilde{Q}$, modulo all P_i 's in time $O(\mathcal{D}(P, Q) + \mathcal{C}(P, Q))$. Finally, given further the generator (G, H) in $\mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha}$, it follows from Lemma 3 that the polynomials $\gamma_1, \dots, \gamma_\alpha$ and $\eta_1, \dots, \eta_\alpha$ can be obtained in time $O(\mathcal{D}(P) + \alpha \mathcal{C}(P))$ and $O(\mathcal{D}(Q) + \alpha \mathcal{C}(Q))$, respectively. Thus, the overall cost of deducing the needed polynomials from P, Q, G, H is in $O(\mathcal{D}(P, Q) + \alpha \mathcal{C}(P, Q))$.

We now turn to the cost of the matrix-vector products performed when applying Theorem 13 to the evaluation of Au or $A'u$ for some u in $\mathbb{F}^n$. Lemma 9 shows that the cost of multiplying $\mathbb{Y}_Q$ by u is $O(\mathcal{M}(n))$, which is $O(\mathcal{M}(p))$. By Lemma 4, the cost for $\mathbb{X}_Q$ is $O(\mathcal{D}(Q) + \mathcal{C}(Q))$. The inner sum amounts to $O(\alpha)$ multiplications modulo P or Q , for a total of $O(\alpha \mathcal{M}(p))$. Lemma 4 also shows that the cost for $\mathbb{W}_P$ is $O(\mathcal{D}(P) + \mathcal{C}(P))$. Finally, multiplying by $\mathbb{V}_{P, Q}$ or $\mathbb{V}'_{P, Q}$ amounts to performing modular multiplications modulo all P_i 's, and this can be done in time $O(\mathcal{M}(m)) \subset O(\mathcal{M}(p))$. Hence the cost of all these matrix-vector products is in $O(\mathcal{D}(P, Q) + \mathcal{C}(P, Q) + \alpha \mathcal{M}(p))$; this fits the announced bound, since $\mathcal{M}(p) \leq \mathcal{C}(P, Q)$. $\square$

The rest of this section is devoted to proving Theorem 13 above.

3.1. Sylvester operator $\nabla_{\mathbb{M}_P, \mathbb{M}_Q^t}$. Assuming that $\gcd(P, Q) = 1$, we show how to reconstruct the unique matrix $A \in \mathbb{F}^{m \times n}$ such that $\nabla_{\mathbb{M}_P, \mathbb{M}_Q^t}(A) = GH^t$. We first show how to reconstruct all blocks in an ad-hoc block decomposition of A , then use Chinese remaindering.

Step 1: reconstruction of the blocks of A . Partitioning the matrix A into blocks $A_{i,j}$ conformally with the block-diagonal structure of the matrices $\mathbb{M}_P$ and $\mathbb{M}_Q$, we have $\nabla_{\mathbb{M}_{P_i}, \mathbb{M}_{Q_j}^t}(A_{i,j}) = G_i H_j^t$ for all i, j . In this paragraph, we prove a reconstruction formula for each block $A_{i,j}$.

LEMMA 15. *For all $i \leq d$ and $j \leq e$, we have*

$$A_{i,j} = \sum_{k \leq \alpha} \mathbb{M}_{g_{i,k}, P_i} \mathbb{M}_{Q_j^{-1}, P_i, n_j} \mathbb{M}_{h_{j,k}, Q_j} \mathbb{Y}_{Q_j}.$$

Proof. Fix i and j . Note that $\mathbb{M}_{P_i}$ and $\mathbb{M}_{Q_j}$ cannot be simultaneously singular, since P_i and Q_j are coprime. Then, for all $\ell \geq 1$, a slight variation of [35, Theorem 4.8] (with both cases “ $\mathbb{M}_{P_i}$ nonsingular” and “ $\mathbb{M}_{Q_j}$ nonsingular” covered by a single identity) gives

$$\mathbb{M}_{P_i}^\ell A_{i,j} - A_{i,j} (\mathbb{M}_{Q_j}^t)^\ell = \sum_{k \leq \alpha} \mathbb{K}(\mathbb{M}_{P_i}, \mathbf{g}_{i,k}, \ell) \mathbb{J}_\ell \mathbb{K}(\mathbb{M}_{Q_j}, \mathbf{h}_{j,k}, \ell)^t,$$

where $\mathbb{K}(\mathbb{M}_{P_i}, \mathbf{g}_{i,k}, \ell) \in \mathbb{F}^{m_i \times \ell}$ denotes the Krylov matrix of column length ℓ associated with the matrix $\mathbb{M}_{P_i}$ and the vector $\mathbf{g}_{i,k}$.

Writing $Q_j = q_{j,0} + \dots + q_{j,n_j} x^{n_j}$, we obtain, for $1 \leq \ell \leq n_j$,

$$\begin{aligned} q_{j,\ell} \mathbb{M}_{P_i}^\ell A_{i,j} - A_{i,j} q_{j,\ell} (\mathbb{M}_{Q_j}^t)^\ell &= \sum_{k \leq \alpha} \mathbb{K}(\mathbb{M}_{P_i}, \mathbf{g}_{i,k}, \ell) q_{j,\ell} \mathbb{J}_\ell \mathbb{K}(\mathbb{M}_{Q_j}, \mathbf{h}_{j,k}, \ell)^t \\ &= \sum_{k \leq \alpha} \mathbb{K}(\mathbb{M}_{P_i}, \mathbf{g}_{i,k}, n_j) q_{j,\ell} \mathbb{J}_{\ell, n_j} \mathbb{K}(\mathbb{M}_{Q_j}, \mathbf{h}_{j,k}, n_j)^t, \end{aligned}$$

since we can inflate all matrices $\mathbb{J}_\ell$ to size $n_j \times n_j$, replacing them by $\mathbb{J}_{\ell, n_j}$. Note that the above equality then also holds for $\ell = 0$, since $\mathbb{J}_{0, n_j} = 0$ by convention. Summing over all $\ell = 0, \dots, n_j$ and using the fact that $Q_j(\mathbb{M}_{Q_j}^t) = Q_j(\mathbb{M}_{Q_j})^t = 0$, we deduce

$$Q_j(\mathbb{M}_{P_i}) A_{i,j} = \sum_{k \leq \alpha} \mathbb{K}(\mathbb{M}_{P_i}, \mathbf{g}_{i,k}, n_j) \mathbb{Y}_{Q_j} \mathbb{K}(\mathbb{M}_{Q_j}, \mathbf{h}_{j,k}, n_j)^t.$$

Using the first part of Lemma 8 together with the fact that for $\gcd(P_i, Q_j) = 1$ the matrix $Q_j(\mathbb{M}_{P_i}) = \mathbb{M}_{Q_j, P_i}$ is invertible with inverse $\mathbb{M}_{Q_j^{-1}, P_i}^{-1} = \mathbb{M}_{Q_j^{-1}, P_i}$, we obtain

$$A_{i,j} = \sum_{k \leq \alpha} \mathbb{M}_{Q_j^{-1}, P_i} \mathbb{M}_{g_{i,k}, P_i, n_j} \mathbb{Y}_{Q_j} \mathbb{M}_{h_{j,k}, Q_j}^t.$$

Hence, applying Lemma 7 to $h_{j,k}$ and since multiplication matrices mod P_i commute,

$$A_{i,j} = \sum_{k \leq \alpha} \mathbb{M}_{g_{i,k}, P_i} \mathbb{M}_{Q_j^{-1}, P_i, n_j} \mathbb{M}_{h_{j,k}, Q_j} \mathbb{Y}_{Q_j}.$$

□

Step 2: a first reconstruction formula for A. Let $\mathbb{B}_{\mathbf{P}, \mathbf{Q}}$ be the $m \times n$ block matrix

$$\mathbb{B}_{\mathbf{P}, \mathbf{Q}} = \left[\mathbb{M}_{Q_j^{-1}, P_i, n_j} \right]_{\substack{1 \leq i \leq d \\ 1 \leq j \leq e}}.$$

Similarly, for $k \leq \alpha$, we define $\mathbf{C}_k$ and $\mathbf{D}_k$ as the block diagonal matrices, with respectively the multiplication matrices $(\mathbb{M}_{g_{i,k}, P_i})_{i \leq d}$ and $(\mathbb{M}_{h_{j,k}, Q_j})_{j \leq e}$ on the diagonal.

Putting all blocks $\mathbf{A}_{i,j}$ together, we deduce the following reconstruction formula for $\mathbf{A}$; the proof is a straightforward application of the previous lemma.

LEMMA 16. *We have*

$$\mathbf{A} = \left(\sum_{k \leq \alpha} \mathbf{C}_k \mathbb{B}_{\mathbf{P}, \mathbf{Q}} \mathbf{D}_k \right) \mathbb{Y}_{\mathbf{Q}}.$$

Step 3: using a factorization of $\mathbb{B}_{\mathbf{P}, \mathbf{Q}}$. Next, we use the fact that $\mathbb{B}_{\mathbf{P}, \mathbf{Q}}$ has a block structure similar to a Cauchy matrix to factor it explicitly. For this, we introduce a matrix related to $\mathbb{W}_{\mathbf{P}}$: we let $\mathbb{W}_{\mathbf{P}, n}$ be the $m \times n$ matrix of multiple reduction modulo $P_1, \dots, P_d$, taking as input a polynomial of degree less than n , instead of m for $\mathbb{W}_{\mathbf{P}}$.

LEMMA 17. *The equality $\mathbb{B}_{\mathbf{P}, \mathbf{Q}} = \mathbb{V}_{\mathbf{P}, \mathbf{Q}} \mathbb{W}_{\mathbf{P}, n} \mathbb{X}_{\mathbf{Q}}$ holds, where $\mathbb{V}_{\mathbf{P}, \mathbf{Q}} \in \mathbb{F}^{m \times m}$ is the block-diagonal matrix with blocks $\mathbb{M}_{Q^{-1}, P_i}$, for $i = 1, \dots, d$.*

Proof. Observe that, given the coefficient vectors of polynomials $F_1, \dots, F_e$, with $F_j \in \mathbb{F}[x]_{n_j}$ for all j , the matrix $\mathbb{B}_{\mathbf{P}, \mathbf{Q}}$ returns the coefficients of

$$\begin{aligned} G_i &= \frac{F_1}{Q_1} + \dots + \frac{F_e}{Q_e} \bmod P_i \\ &= \frac{F_1 Q_2 \dots Q_e + \dots + Q_1 \dots Q_{e-1} F_e}{Q_1 \dots Q_e} \bmod P_i, \end{aligned}$$

for $i = 1, \dots, d$. Computing the polynomials G_i can be done as follows:

- compute $H = F_1 Q_2 \dots Q_e + \dots + Q_1 \dots Q_{e-1} F_e$ by calling `combP`;
- compute $H_i = H \bmod P_i$, for $i = 1, \dots, d$;
- deduce $G_i = H_i / Q \bmod P_i$, for $i = 1, \dots, d$.

This proves the requested factorization of $\mathbb{B}_{\mathbf{P}, \mathbf{Q}}$. □

As a result, we deduce the following equalities for $\mathbf{A}$:

$$\begin{aligned} \mathbf{A} &= \left(\sum_{k \leq \alpha} \mathbf{C}_k \mathbb{V}_{\mathbf{P}, \mathbf{Q}} \mathbb{W}_{\mathbf{P}, n} \mathbb{X}_{\mathbf{Q}} \mathbf{D}_k \right) \mathbb{Y}_{\mathbf{Q}} \\ &= \mathbb{V}_{\mathbf{P}, \mathbf{Q}} \left(\sum_{k \leq \alpha} \mathbf{C}_k \mathbb{W}_{\mathbf{P}, n} \mathbb{X}_{\mathbf{Q}} \mathbf{D}_k \right) \mathbb{Y}_{\mathbf{Q}}. \end{aligned}$$

The last equality is due to the fact that the block-diagonal matrices $\mathbf{C}_k$ and $\mathbb{V}_{\mathbf{P}, \mathbf{Q}}$ commute (since all pairwise corresponding blocks are multiplication matrices modulo the same polynomials P_i).

Step 4: using Chinese remaindering. The next step will allow us to take $\mathbb{W}_{\mathbf{P}, n}$ and $\mathbb{X}_{\mathbf{P}}$ out of the inner sum. For any polynomial H , any $i \leq d$ and any $k \leq \alpha$, it is equivalent to (i) reduce H modulo P_i and multiply it by $g_{i,k}$ modulo P_i and (ii) multiply H by the polynomial γ_k (defined above) modulo P , and reduce it modulo P_i . In other words, we have the commutation relation $\mathbf{C}_k \mathbb{W}_{\mathbf{P}, n} = \mathbb{W}_{\mathbf{P}} \mathbb{M}_{\gamma_k, P, n}$. Similarly, $\mathbb{X}_{\mathbf{Q}} \mathbf{D}_k = \mathbb{M}_{\eta_k, Q} \mathbb{X}_{\mathbf{Q}}$ and this concludes the proof of the first part of Theorem 13.

3.2. Stein operator $\Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$. The proof for Stein operator case follows the same steps as for the Sylvester case. Let A' be such that $\Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}(A') = \mathbf{G}\mathbf{H}^t$. Just like we decomposed $\mathbf{A}$ into blocks $\mathbf{A}_{i,j}$, we decompose A' into blocks $A'_{i,j}$, such that $\Delta_{\mathbb{M}_{P_i}, \mathbb{M}_{Q_j}^t}(A'_{i,j}) = \mathbf{G}_i \mathbf{H}_j^t$. Extending the notation used above, we write $\widetilde{Q}_j = \text{rev}(Q_j, n_j)$, so that we have $\widetilde{Q} = \widetilde{Q}_1 \cdots \widetilde{Q}_e$. The following lemma is then the analogue of Lemma 15 for Stein operators and, as detailed in Appendix B, can be proved in the same way using [35, Theorem 4.7].

LEMMA 18. *For all $i \leq d$ and $j \leq e$, we have*

$$A'_{i,j} = \sum_{k \leq \alpha} \mathbb{M}_{g_{i,k}, P_i} \mathbb{M}_{\widetilde{Q}_j^{-1}, P_i, n_j} \mathbb{J}_{n_j} \mathbb{M}_{h_{j,k}, Q_j} \mathbb{Y}_{Q_j}.$$

Mimicking the construction in the previous section, we introduce the $m \times n$ block matrix $\mathbb{B}'_{\mathbf{P}, \mathbf{Q}}$ given by

$$\mathbb{B}'_{\mathbf{P}, \mathbf{Q}} = \left[\mathbb{M}_{\widetilde{Q}_j^{-1}, P_i, n_j} \right]_{\substack{1 \leq i \leq d \\ 1 \leq j \leq e}}.$$

We will use again the matrices $\mathbf{C}_k$ and $\mathbf{D}_k$ introduced before, as well as the block-diagonal matrix $\mathbb{D}(\mathbb{J}_{n_j})$ having blocks $\mathbb{J}_{n_j}$ on the diagonal. This leads us to the following analogue of Lemma 16, whose proof is straightforward.

LEMMA 19. *We have*

$$A' = \left(\sum_{k \leq \alpha} \mathbf{C}_k \mathbb{B}'_{\mathbf{P}, \mathbf{Q}} \mathbb{D}(\mathbb{J}_{n_j}) \mathbf{D}_k \right) \mathbb{Y}_{\mathbf{Q}}.$$

The next step is to use the following factorization of $\mathbb{B}'_{\mathbf{P}, \mathbf{Q}}$, or more precisely of $\mathbb{B}'_{\mathbf{P}, \mathbf{Q}} \mathbb{D}(\mathbb{J}_{n_j})$. The proof is the same as that of Lemma 17, up to taking into account the reversals induced by the matrices $\mathbb{J}_{n_j}$.

LEMMA 20. *The equality $\mathbb{B}'_{\mathbf{P}, \mathbf{Q}} \mathbb{D}(\mathbb{J}_{n_j}) = \mathbb{V}'_{\mathbf{P}, \mathbf{Q}} \mathbb{W}_{\mathbf{P}, n} \mathbb{J}_n \mathbb{X}_{\mathbf{Q}}$ holds, where $\mathbb{V}'_{\mathbf{P}, \mathbf{Q}} \in \mathbb{F}^{m \times m}$ is the block-diagonal matrix with blocks $\mathbb{M}_{\widetilde{Q}_j^{-1}, P_i}$, for $i = 1, \dots, d$.*

We conclude the proof of our theorem as before, using the relations

$$\begin{aligned} \mathbf{C}_k \mathbb{V}'_{\mathbf{P}, \mathbf{Q}} &= \mathbb{V}'_{\mathbf{P}, \mathbf{Q}} \mathbf{C}_k, \\ \mathbf{C}_k \mathbb{W}_{\mathbf{P}, n} &= \mathbb{W}_{\mathbf{P}} \mathbb{M}_{\gamma_k, P, n}, \\ \mathbb{X}_{\mathbf{Q}} \mathbf{D}_k &= \mathbb{M}_{\eta_k, Q} \mathbb{X}_{\mathbf{Q}}. \end{aligned}$$

4. Using operator equivalences. Let $\mathbf{P}, \mathbf{Q}$ be as in Theorems 1 and 2, and let as before $p = \max(m, n)$. We are now going to extend the complexity estimates for matrix-vector multiplication given in Corollary 14 to more operators (not only the basic ones), by providing reductions to the Hankel operator $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$.

THEOREM 21. *Suppose that $\mathcal{H}_{\mathbf{P}}$ and $\mathcal{H}_{\mathbf{Q}}$ hold. Then for any displacement operator $\mathcal{L}$ associated with $(\mathbf{P}, \mathbf{Q})$, we can take*

$$\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) \leq \mathcal{C}_{\text{mul}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}, \alpha + 2, \beta) + O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + (\alpha + \beta)\mathcal{C}(\mathbf{P}, \mathbf{Q})),$$

$$\mathcal{C}_{\text{solve}}(\mathcal{L}, \alpha) \leq \mathcal{C}_{\text{solve}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}, \alpha + 2) + O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha\mathcal{C}(\mathbf{P}, \mathbf{Q})),$$

$$\mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha) \leq \mathcal{C}_{\text{inv}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{m,1}^t}, \alpha + 2) + O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha\mathcal{C}(\mathbf{P}, \mathbf{Q}) + \alpha^{\omega-1}m).$$

The rest of this section is devoted to the proof of this theorem. In what follows, we write several formulas involving some matrices G and H . In these formulas, G and H will always be taken in $\mathbb{F}^{m \times \alpha}$ and $\mathbb{F}^{n \times \alpha}$, respectively, for some $\alpha \geq 1$ (and with $n = m$ when dealing with matrix inversion).

4.1. Reduction to basic operators. Table 4 shows that if a matrix A is structured for one of the operators associated with (P, Q) , then simple pre- and post-multiplications transform A into a matrix which is structured for one of the two *basic* operators associated with (P, Q) . All equivalences in this table follow directly from the identities $Y_P M_P^t = M_P Y_P$ and $Y_Q M_Q^t = M_Q Y_Q$ (Lemma 10) and from Y_P and Y_Q being invertible and symmetric. Combining these equivalences with Lemma 9 will lead to the cost bounds in Lemma 22 below, expressed in terms of basic operators.

TABLE 4
Reduction to the basic operators ∇_{M_P, M_Q^t} and Δ_{M_P, M_Q^t} .

$$\begin{aligned} \nabla_{M_P, M_Q}(A) &= G H^t \iff \nabla_{M_P, M_Q^t}(A Y_Q) = G (Y_Q H)^t \\ \nabla_{M_P^t, M_Q}(A) &= G H^t \iff \nabla_{M_P, M_Q^t}(Y_P A) = (Y_P G) H^t \\ \nabla_{M_P^t, M_Q}(A) &= G H^t \iff \nabla_{M_P, M_Q^t}(Y_P A Y_Q) = (Y_P G)(Y_Q H)^t \\ \Delta_{M_P, M_Q}(A) &= G H^t \iff \Delta_{M_P, M_Q^t}(A Y_Q) = G (Y_Q H)^t \\ \Delta_{M_P^t, M_Q}(A) &= G H^t \iff \Delta_{M_P, M_Q^t}(Y_P A) = (Y_P G) H^t \\ \Delta_{M_P^t, M_Q}(A) &= G H^t \iff \Delta_{M_P, M_Q^t}(Y_P A Y_Q) = (Y_P G)(Y_Q H)^t. \end{aligned}$$

LEMMA 22. Let $\mathcal{L}$ be a displacement operator associated with (P, Q) . Then:

- If $\mathcal{L}$ is a Sylvester operator, we can take

$$\begin{aligned} \mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) &\leq \mathcal{C}_{\text{mul}}(\nabla_{M_P, M_Q^t}, \alpha, \beta) + O((\alpha + \beta)\mathcal{M}(p)), \\ \mathcal{C}_{\text{solve}}(\mathcal{L}, \alpha) &\leq \mathcal{C}_{\text{solve}}(\nabla_{M_P, M_Q^t}, \alpha) + O(\alpha\mathcal{M}(p)), \\ \mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha) &\leq \mathcal{C}_{\text{inv}}(\nabla_{M_P, M_Q^t}, \alpha) + O(\alpha\mathcal{M}(m)); \end{aligned}$$

- If $\mathcal{L}$ is a Stein operator, we can take

$$\begin{aligned} \mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) &\leq \mathcal{C}_{\text{mul}}(\Delta_{M_P, M_Q^t}, \alpha, \beta) + O((\alpha + \beta)\mathcal{M}(p)), \\ \mathcal{C}_{\text{solve}}(\mathcal{L}, \alpha) &\leq \mathcal{C}_{\text{solve}}(\Delta_{M_P, M_Q^t}, \alpha) + O(\alpha\mathcal{M}(p)), \\ \mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha) &\leq \mathcal{C}_{\text{inv}}(\Delta_{M_P, M_Q^t}, \alpha) + O(\alpha\mathcal{M}(m)). \end{aligned}$$

Proof. We give the proof for the Sylvester operator $\mathcal{L} = \nabla_{M_P^t, M_Q}$, which appears in the third row of Table 4; the other five cases in that table can be handled similarly.

Suppose first that, given an $\mathcal{L}$ -generator (G, H) for $A \in \mathbb{F}^{m \times n}$ as well as a matrix $B \in \mathbb{F}^{n \times \beta}$, we want to compute AB . We begin by computing a $\nabla_{M_P, M_Q^t}(A')$ -generator $(Y_P G, Y_Q H)$ for $A' = Y_P A Y_Q$ (cf. the third row of Table 4). Lemma 9 implies that this takes time $O(\alpha\mathcal{M}(p))$. Then, we evaluate $AB = Y_P^{-1} A' Y_Q^{-1} B$ from right to left. The products by Y_P^{-1} and Y_Q^{-1} take time $O(\beta\mathcal{M}(p))$, and the product by A' takes time $\mathcal{C}_{\text{mul}}(\nabla_{M_P, M_Q^t}, \alpha, \beta)$; thus, the first claim is proved.

Suppose now that we are given a vector $\mathbf{b} \in \mathbb{K}^m$ in addition to the matrix $\mathbf{A}$. In order to solve the system $\mathbf{A}\mathbf{x} = \mathbf{b}$, we solve $\mathbf{A}'\mathbf{x}' = \mathbf{b}'$, with $\mathbf{A}'$ as defined above and $\mathbf{b}' = \mathbb{Y}_{\mathbf{P}}\mathbf{b}$. The latter system admits a solution if and only if the former one does; if $\mathbf{x}'$ is a solution of $\mathbf{A}'\mathbf{x}' = \mathbf{b}'$, then $\mathbf{x} = \mathbb{Y}_{\mathbf{Q}}\mathbf{x}'$ is a solution of $\mathbf{A}\mathbf{x} = \mathbf{b}$. Hence, as before, we set up an $\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ -generator $(\mathbb{Y}_{\mathbf{P}}\mathbf{G}, \mathbb{Y}_{\mathbf{Q}}\mathbf{H})$ for $\mathbf{A}'$, using $O(\alpha\mathcal{M}(p))$ operations in $\mathbb{F}$, and compute $\mathbf{b}'$ in time $O(\mathcal{M}(m))$. Then, in time $\mathcal{C}_{\text{solve}}(\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}, \alpha)$ we either assert that the system $\mathbf{A}'\mathbf{x}' = \mathbf{b}'$ has no solution, or find such a solution $\mathbf{x}'$. Finally, we recover $\mathbf{x}$ in time $O(\mathcal{M}(n))$. This proves the second claim.

Finally, assume that $m = n$ and consider the question of inverting $\mathbf{A}$. This matrix is invertible if and only if the matrix $\mathbf{A}'$ defined above is invertible. Again, we set up a $\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ -generator $(\mathbb{Y}_{\mathbf{P}}\mathbf{G}, \mathbb{Y}_{\mathbf{Q}}\mathbf{H})$ for $\mathbf{A}'$, using $O(\alpha\mathcal{M}(m))$ operations in $\mathbb{F}$. Then, in time $\mathcal{C}_{\text{inv}}(\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}, \alpha)$ we either assert that $\mathbf{A}'$ is not invertible or deduce a $\nabla_{\mathbb{M}_{\mathbf{Q}}, \mathbb{M}_{\mathbf{P}}}$ -generator for $\mathbf{A}'^{-1}$, say $(\mathbf{G}'_{\text{inv}}, \mathbf{H}'_{\text{inv}})$. Finally, if $\mathbf{A}'$ is invertible then, using $\mathbb{M}_{\mathbf{P}}^t = \mathbb{Y}_{\mathbf{P}}^{-1}\mathbb{M}_{\mathbf{P}}\mathbb{Y}_{\mathbf{P}}$ and $\mathbb{M}_{\mathbf{Q}} = \mathbb{Y}_{\mathbf{Q}}\mathbb{M}_{\mathbf{Q}}^t\mathbb{Y}_{\mathbf{Q}}^{-1}$, we obtain a $\nabla_{\mathbb{M}_{\mathbf{Q}}, \mathbb{M}_{\mathbf{P}}}$ -generator $(\mathbb{Y}_{\mathbf{Q}}\mathbf{G}'_{\text{inv}}, \mathbb{Y}_{\mathbf{P}}\mathbf{H}'_{\text{inv}})$ for $\mathbf{A}^{-1}$ in time $O(\alpha\mathcal{M}(m))$. This proves the last claim. $\square$

4.2. Reduction to the Hankel case. Our second reduction is less straightforward: we use Pan's idea of multiplicative transformation of operators [30] to reduce *basic* operators to an operator of Hankel type. There is some flexibility in the choice of the target operator, here the Sylvester operator $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$; the only (natural) requirement is that this target operator remains invertible.

The following proposition summarizes the transformation process. Although the formulas are long, they describe simple processes: for an operation such as multiplication, inversion or system solving, this amounts to e.g. the analogue operation for an operator of Hankel type, several products with simple matrices derived from $\mathbf{P}$ and $\mathbf{Q}$, and $O(1)$ matrix-vector products with the input matrix or its transpose.

PROPOSITION 23. *Let $\mathcal{L} \in \{\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}, \Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}\}$, and suppose that $\mathcal{H}_{\mathbf{P}}$ and $\mathcal{H}_{\mathbf{Q}}$ hold. Then*

$$\begin{aligned} (5) \quad \mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) &\leq \mathcal{C}_{\text{mul}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}, \alpha + 2, \beta) + O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + (\alpha + \beta)\mathcal{C}(\mathbf{P}, \mathbf{Q})), \\ (6) \quad \mathcal{C}_{\text{solve}}(\mathcal{L}, \alpha) &\leq \mathcal{C}_{\text{solve}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}, \alpha + 2) + O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha\mathcal{C}(\mathbf{P}, \mathbf{Q})), \\ (7) \quad \mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha) &\leq \mathcal{C}_{\text{inv}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}, \alpha + 2) + O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha\mathcal{C}(\mathbf{P}, \mathbf{Q}) + \alpha^{\omega-1}m). \end{aligned}$$

The proof of Proposition 23 will occupy the rest of this subsection. Before that, though, we point out that Theorem 21 follows directly from this proposition and Lemma 22. In particular, since $p \leq \mathcal{M}(p) \leq \mathcal{M}(m) + \mathcal{M}(n) \leq \mathcal{C}(\mathbf{P}, \mathbf{Q})$ for $p = \max(m, n)$, overheads such as $O(\alpha\mathcal{M}(m))$ or $O(\alpha\mathcal{M}(p))$ that appear in that lemma are absorbed into terms such as $O(\alpha\mathcal{C}(\mathbf{P}, \mathbf{Q}))$ that appear in the proposition.

Preliminaries. To establish each of the claims in the proposition above, it will be useful to have simple matrices $\mathbf{L}$ and $\mathbf{R}$ for pre-/post-multiplying $\mathbf{A}$ and whose displacement rank with respect to $\nabla_{\mathbb{Z}_{m,0}, \mathbb{M}_{\mathbf{P}}}$ and $\nabla_{\mathbb{M}_{\mathbf{Q}}, \mathbb{Z}_{n,1}^t}$, respectively, is small. Lemma 24 below shows that a possible choice, leading to displacement ranks at most 1, is

$$\mathbf{L} = \mathbb{J}_m \mathbb{W}_{\mathbf{P}}^t \mathbb{Y}_{\mathbf{P}}^{-1} \quad \text{and} \quad \mathbf{R} = \mathbb{Y}_{\mathbf{Q}}^{-1} \mathbb{W}_{\mathbf{Q}} \mathbb{J}_n.$$

In order to define generators for such matrices, we first rewrite the products $P = P_1 \cdots P_d$ and $Q = Q_1 \cdots Q_e$ as $P = p_0 + \cdots + p_{m-1}x^{m-1} + x^m$ and $Q = q_0 + \cdots + q_{n-1}x^{n-1} + x^n$, and let $\mathbf{m} = [p_0 \cdots p_{m-1}]^t$ and $\mathbf{n} = [q_0 \cdots q_{n-1}]^t$ be the coefficient

vectors of the polynomials $P - x^m$ and $Q - x^n$. Then, we let

$$\mathbf{t} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \in \mathbb{F}^m, \quad \mathbf{u} = \mathbb{Y}_{\mathbf{P}}^{-1} \mathbb{W}_{\mathbf{P}} \mathbf{m} \in \mathbb{F}^m, \quad \mathbf{s} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \in \mathbb{F}^n, \quad \mathbf{r} = -\mathbb{Y}_{\mathbf{Q}}^{-1} \mathbb{W}_{\mathbf{Q}} (\mathbf{n} + \mathbf{s}) \in \mathbb{F}^n.$$

LEMMA 24. *The relations $\nabla_{\mathbb{Z}_{m,0}, \mathbb{M}_{\mathbf{P}}}(\mathbf{L}) = \mathbf{t} \mathbf{u}^t$ and $\nabla_{\mathbb{M}_{\mathbf{Q}}, \mathbb{Z}_{n,1}^t}(\mathbf{R}) = \mathbf{r} \mathbf{s}^t$ hold.*

Proof. According to the second part of Lemma 11, $\mathbb{Y}_{\mathbf{P}}^{-1} \mathbb{W}_{\mathbf{P}}$ is equal to the Krylov matrix $\mathbb{K}(\mathbb{M}_{\mathbf{P}}^t, \mathbf{w}, m)$, so that the columns of $\mathbf{L}^t = \mathbb{Y}_{\mathbf{P}}^{-1} \mathbb{W}_{\mathbf{P}} \mathbb{J}_m$ are $(\mathbb{M}_{\mathbf{P}}^t)^{m-1} \mathbf{w}, \dots, \mathbb{M}_{\mathbf{P}}^t \mathbf{w}, \mathbf{w}$. Hence only the first column of $\mathbb{M}_{\mathbf{P}}^t \mathbf{L}^t - \mathbf{L}^t \mathbb{Z}_{m,0}^t$ is nonzero, and it is equal to $\tilde{\mathbf{u}} := (\mathbb{M}_{\mathbf{P}}^t)^m \mathbf{w}$. After transposition, this leads to $\nabla_{\mathbb{Z}_{m,0}, \mathbb{M}_{\mathbf{P}}}(\mathbf{L}) = -\mathbf{t} \tilde{\mathbf{u}}^t$. We can now check that $\tilde{\mathbf{u}} = -\mathbf{u}$ as follows. First, from Lemma 10 and since $\mathbb{Y}_{\mathbf{P}}$ is invertible, we see that $\tilde{\mathbf{u}} = \mathbb{Y}_{\mathbf{P}}^{-1} \mathbb{M}_{\mathbf{P}}^m \mathbb{Y}_{\mathbf{P}} \mathbf{w}$. Then, the shape of $\mathbf{w}$ implies that $\mathbb{Y}_{\mathbf{P}} \mathbf{w}$ is the vector $\mathbf{v}$ defined in the first part of Lemma 11, so that $\tilde{\mathbf{u}} = \mathbb{Y}_{\mathbf{P}}^{-1} \mathbb{M}_{\mathbf{P}}^m \mathbf{v}$. Now, the special shape of $\mathbf{v}$ implies that $\mathbb{M}_{\mathbf{P}}^m \mathbf{v}$ is the vector whose i th subvector of length m_i contains the coefficients of $x^m \bmod P_i = -(P - x^m) \bmod P_i$. In other words, $\mathbb{M}_{\mathbf{P}}^m \mathbf{v} = -\mathbb{W}_{\mathbf{P}} \mathbf{m}$. This shows that $\tilde{\mathbf{u}} = -\mathbb{Y}_{\mathbf{P}}^{-1} \mathbb{W}_{\mathbf{P}} \mathbf{m} = -\mathbf{u}$, which concludes the proof of the first relation.

For the second relation, the proof is the same, taking into account that we are now considering the operator $\nabla_{\mathbb{M}_{\mathbf{Q}}, \mathbb{Z}_{n,1}^t}$ and that $\mathbb{Z}_{n,1}^t = \mathbb{Z}_{n,0}^t + \mathbb{J}_n \mathbf{s} \mathbf{s}^t$. $\square$

Proof of (5) for $\mathcal{L} = \nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$. Let $\mathbf{A}$ be in $\mathbb{F}^{m \times n}$ and let $(\mathbf{G}, \mathbf{H})$ be a $\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ -generator for $\mathbf{A}$. With $\mathbf{L}, \mathbf{R}, \mathbf{r}, \mathbf{s}, \mathbf{t}, \mathbf{u}$ as in the previous paragraph, define further

$$\mathbf{G}' = [\mathbf{t} \mid \mathbf{L} \mathbf{G} \mid \mathbf{L} \mathbf{A} \mathbf{r}] \quad \text{and} \quad \mathbf{H}' = [\mathbf{R}^t \mathbf{A}^t \mathbf{u} \mid \mathbf{R}^t \mathbf{H} \mid \mathbf{s}].$$

LEMMA 25. *The matrix $\mathbf{A}' = \mathbf{L} \mathbf{A} \mathbf{R}$ satisfies $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}(\mathbf{A}') = \mathbf{G}' \mathbf{H}'^t$.*

Proof. Applying the general formula

$$(8) \quad \nabla_{\mathbf{M}, \mathbf{N}}(\mathbf{A} \mathbf{B} \mathbf{C}) = \nabla_{\mathbf{M}, \mathbf{Q}}(\mathbf{A}) \mathbf{B} \mathbf{C} + \mathbf{A} \nabla_{\mathbf{Q}, \mathbf{R}}(\mathbf{B}) \mathbf{C} + \mathbf{A} \mathbf{B} \nabla_{\mathbf{R}, \mathbf{N}}(\mathbf{C})$$

(which follows directly from [33, Theorem 1.5.4]) and then using Lemma 24, we obtain

$$\begin{aligned} \nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}(\mathbf{L} \mathbf{A} \mathbf{R}) &= \nabla_{\mathbb{Z}_{m,0}, \mathbb{M}_{\mathbf{P}}}(\mathbf{L}) \mathbf{A} \mathbf{R} + \mathbf{L} \nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}(\mathbf{A}) \mathbf{R} + \mathbf{L} \mathbf{A} \nabla_{\mathbb{M}_{\mathbf{Q}}, \mathbb{Z}_{n,1}^t}(\mathbf{R}) \\ &= \mathbf{t} \mathbf{u}^t \mathbf{A} \mathbf{R} + \mathbf{L} \mathbf{G} \mathbf{H}^t \mathbf{R} + \mathbf{L} \mathbf{A} \mathbf{r} \mathbf{s}^t, \end{aligned}$$

which is the announced equality. $\square$

To compute a product of the form $\mathbf{A} \mathbf{B}$ with $\mathbf{B} \in \mathbb{F}^{n \times \beta}$, we first compute the matrices $\mathbf{G}'$ and $\mathbf{H}'$ described above. To obtain $\mathbf{G}'$, it suffices to set up the vectors $\mathbf{r}$ and $\mathbf{A} \mathbf{r}$ and to compute $\alpha + 1$ matrix-vector products by $\mathbf{L}$. Given $(\mathbf{G}, \mathbf{H})$ and $\mathbf{r}$, Corollary 14 implies that $\mathbf{A} \mathbf{r}$ is obtained in time

$$O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha \mathcal{C}(\mathbf{P}, \mathbf{Q})).$$

On the other hand, it follows from Lemmas 4 and 9 that the vector $\mathbf{r}$ is obtained in time $O(\mathcal{C}(\mathbf{Q}) + \mathcal{D}(\mathbf{Q}))$ and, after some precomputation of time $O(\mathcal{C}(\mathbf{P}) + \mathcal{D}(\mathbf{P}))$, that the $\alpha + 1$ products by $\mathbf{L}$ are obtained in time $O(\alpha \mathcal{C}(\mathbf{P}))$. Thus, overall, $\mathbf{G}'$ is obtained in time $O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha \mathcal{C}(\mathbf{P}, \mathbf{Q}))$.

We proceed similarly for H' . The only differences are that we have to do matrix-vector products involving R^t and A^t instead of L and A . For the former, Lemmas 4 and 9 show that after a precomputation of cost $O(\mathcal{C}(\mathbf{Q}) + \mathcal{D}(\mathbf{Q}))$, the cost of one such product is $O(\mathcal{C}(\mathbf{Q}))$. For the latter, recall that, as pointed out in the introduction, from the given $\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ -generator of A , we can deduce in negligible time a $\nabla_{\mathbb{M}_{\mathbf{Q}}, \mathbb{M}_{\mathbf{P}}^t}$ -generator of A^t of the same length α . This allows us to do matrix-vector products with A^t with the same asymptotic cost

$$O(\mathcal{D}(\mathbf{Q}, \mathbf{P}) + \alpha \mathcal{C}(\mathbf{Q}, \mathbf{P})) = O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha \mathcal{C}(\mathbf{P}, \mathbf{Q}))$$

as for A . Thus, the overall cost for computing H' is asymptotically the same as for G' .

Let us now bound the cost of deducing the product AB from G', H', B . Note first that under the coprimality assumptions $\mathcal{H}_{\mathbf{P}}$ and $\mathcal{H}_{\mathbf{Q}}$ the matrices $\mathbb{W}_{\mathbf{P}}$ and $\mathbb{W}_{\mathbf{Q}}$ are invertible, and so are L and R . Consequently, $AB = L^{-1}A'R^{-1}B$ and it suffices to bound the cost of each of the three products $B' := R^{-1}B$, $B'' := A'B'$, and $L^{-1}B''$. By reusing the same precomputation as for H' and applying again Lemmas 4 and 9, we obtain B' for an additional cost of $O(\beta \mathcal{C}(\mathbf{Q}))$, via β matrix-vector products by R^{-1} . Then, Lemma 25 says that A' is a Hankel-like matrix for which a $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$ -generator of length $\alpha + 2$ is (G', H') . Hence the cost for deducing B'' from G', H', B' is at most $\mathcal{C}_{\text{mul}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}, \alpha + 2, \beta)$. Finally, $L^{-1}B''$ is obtained for an additional cost of $O(\beta \mathcal{C}(\mathbf{P}))$, by reusing the same precomputation as for G' and by performing β matrix-vector products by L^{-1} .

To summarize, we have shown that (G', H') can be obtained in time $O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha \mathcal{C}(\mathbf{P}, \mathbf{Q}))$ and that AB can be deduced from G', H', B in time $\mathcal{C}_{\text{mul}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}, \alpha + 2, \beta) + O(\beta \mathcal{C}(\mathbf{P}, \mathbf{Q}))$. Adding these two costs thus proves the bound (5) in Proposition 23 for $\mathcal{L} = \nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$.

Proof of (6) for $\mathcal{L} = \nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$. For $A \in \mathbb{F}^{m \times n}$ given by some $\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ -generator (G, H) of length α and given $\mathbf{b}$ in $\mathbb{F}^m$, we now want to find a (nontrivial) solution of $A\mathbf{x} = \mathbf{b}$, or determine that no solution exists.

Define $\mathbf{b}' = L\mathbf{b}$. Because L and R are invertible matrices, solving the system $A'\mathbf{x}' = \mathbf{b}'$ is equivalent to solving $A\mathbf{x} = \mathbf{b}$, with then $\mathbf{x} = R\mathbf{x}'$. As in the previous paragraph, we can compute a generator (G', H') of A' for the operator $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$ in time $O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha \mathcal{C}(\mathbf{P}, \mathbf{Q}))$. The cost of computing $\mathbf{b}'$ from $\mathbf{b}$ (and $\mathbf{x}$ from $\mathbf{x}'$) fits into the same bound, and solving the new system $A'\mathbf{x}' = \mathbf{b}'$ takes time $\mathcal{C}_{\text{solve}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}, \alpha + 2)$. Summing these costs, we prove the second item in the proposition.

Proof of (7) for $\mathcal{L} = \nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$. Assume now that $m = n$ and that $A \in \mathbb{F}^{m \times m}$ is given by a $\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ -generator (G, H) of length α . As before, L and R are invertible by assumption, so that A is invertible if and only if $A' = LAR$ is invertible. By Lemma 25 this matrix A' has displacement rank at most $\alpha + 2$ for $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{m,1}^t}$. Thus, as explained in the introduction, if A' is invertible then its inverse has displacement rank at most $\alpha + 2$ for $\nabla_{\mathbb{Z}_{m,1}^t, \mathbb{Z}_{m,0}}$. The next lemma shows that if $(G'_{\text{inv}}, H'_{\text{inv}})$ is a $\nabla_{\mathbb{Z}_{m,1}^t, \mathbb{Z}_{m,0}}$ -generator for the inverse of A' , then a $\nabla_{\mathbb{M}_{\mathbf{Q}}, \mathbb{M}_{\mathbf{P}}}$ -generator for the inverse of A is given by the matrices

$$G_{\text{inv}} = [r \mid R G'_{\text{inv}} \mid R A'^{-1} t] \quad \text{and} \quad H_{\text{inv}} = [L^t A'^{-t} s \mid L^t H'_{\text{inv}} \mid u].$$

LEMMA 26. *The matrix A^{-1} satisfies $\nabla_{\mathbb{M}_{\mathbf{Q}}, \mathbb{M}_{\mathbf{P}}}^t(A^{-1}) = G_{\text{inv}} H_{\text{inv}}^t$.*

Proof. As for Lemma 25 the proof follows from [33, Theorem 1.5.4] and Lemma 24:

$$\begin{aligned} \nabla_{\mathbb{M}_{\mathbf{Q}}^t, \mathbb{M}_{\mathbf{P}}}(\mathbf{R} \mathbf{A}'^{-1} \mathbf{L}) &= \nabla_{\mathbb{M}_{\mathbf{Q}}^t, \mathbb{Z}_{m,1}^t}(\mathbf{R}) \mathbf{A}'^{-1} \mathbf{L} + \mathbf{R} \nabla_{\mathbb{Z}_{m,1}^t, \mathbb{Z}_{m,0}}(\mathbf{A}'^{-1}) \mathbf{L} \\ &\quad + \mathbf{R} \mathbf{A}'^{-1} \nabla_{\mathbb{Z}_{m,0}, \mathbb{M}_{\mathbf{P}}}(\mathbf{L}) \\ &= \mathbf{r} \mathbf{s}^t \mathbf{A}'^{-1} \mathbf{L} + \mathbf{R} \mathbf{G}'_{\text{inv}} (\mathbf{H}'_{\text{inv}})^t \mathbf{L} + \mathbf{R} \mathbf{A}'^{-1} \mathbf{t} \mathbf{u}^t. \end{aligned}$$

To conclude, we remark that $\mathbf{A}^{-1} = \mathbf{R} \mathbf{A}'^{-1} \mathbf{L}$. $\square$

The bound in (7) can now be established as follows. We begin by computing $(\mathbf{G}', \mathbf{H}')$, which by Lemma 25 is a $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{m,1}^t}$ -generator of length $\alpha + 2$ of $\mathbf{A}'$; as shown in the previous paragraph, this is done in time $O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha \mathcal{C}(\mathbf{P}, \mathbf{Q}))$. Then, in time $\mathcal{C}_{\text{inv}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{m,1}^t}, \alpha + 2)$ we either conclude that $\mathbf{A}'$ is singular, or produce a $\nabla_{\mathbb{Z}_{m,1}^t, \mathbb{Z}_{m,0}}$ -generator $(\mathbf{G}'_{\text{inv}}, \mathbf{H}'_{\text{inv}})$ of length $\alpha + 2$ for the inverse of $\mathbf{A}'$. Finally, if $\mathbf{A}'$ is invertible we proceed in two steps: First, using the same amount of time as for $\mathbf{G}'$ and $\mathbf{H}'$, we set up the matrices $\mathbf{G}_{\text{inv}}$ and $\mathbf{H}_{\text{inv}}$ introduced before Lemma 26. Then we reduce each of these two matrices to arrive at a generator of length α ; this generator compression step can be done in time $O(\alpha^{\omega-1} m)$, in view of [33, Remark 4.6.7]. Overall, these costs add up to the result reported in (7).

Proof of (5)–(7) for $\mathcal{L} = \Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$. We only sketch the proof in these cases. This time, we rely on the (easily verified) general formula

$$\Delta_{\mathbf{M}, \mathbf{N}}(\mathbf{ABC}) = -\nabla_{\mathbf{M}, \mathbf{P}}(\mathbf{A}) \mathbf{B} \mathbf{Q} \mathbf{C} + \mathbf{A} \Delta_{\mathbf{P}, \mathbf{Q}}(\mathbf{B}) \mathbf{C} + \mathbf{M} \mathbf{A} \mathbf{B} \nabla_{\mathbf{Q}, \mathbf{N}}(\mathbf{C}).$$

We use it to write

$$\Delta_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}(\mathbf{L} \mathbf{A} \mathbf{R}) = -\nabla_{\mathbb{Z}_{m,0}, \mathbb{M}_{\mathbf{P}}}(\mathbf{L}) \mathbf{A} \mathbb{M}_{\mathbf{Q}}^t \mathbf{R} + \mathbf{L} \Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}(\mathbf{A}) \mathbf{R} + \mathbb{Z}_{m,0} \mathbf{L} \mathbf{A} \nabla_{\mathbb{M}_{\mathbf{Q}}^t, \mathbb{Z}_{n,1}^t}(\mathbf{R});$$

this allows us to reduce questions (multiplication, linear system solving, inversion) related to $\mathbf{A}$ to the same questions for $\mathbf{A}' = \mathbf{L} \mathbf{A} \mathbf{R}$, where $\mathbf{A}'$ is given through a $\Delta_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$ -generator of length $\alpha + 2$. Then, the relations $\mathbb{Z}_{n,1}^t \mathbb{J}_n = \mathbb{J}_n \mathbb{Z}_{n,1}$ and $\mathbb{Z}_{n,1}^t = \mathbb{Z}_{n,1}^{-1}$ imply

$$\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}(\mathbf{A}' \mathbb{J}_n) = -\Delta_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}(\mathbf{A}') \mathbb{Z}_{n,1} \mathbb{J}_n;$$

this allows us to reduce our problems to computations with the matrix $\mathbf{A}' \mathbb{J}_n$ given by means of a $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$ -generator of length $\alpha + 2$.

For inversion (where $m = n$), inverting $\mathbf{A}' \mathbb{J}_n$ leads to a $\nabla_{\mathbb{Z}_{m,1}^t, \mathbb{Z}_{m,0}}$ -generator of $(\mathbf{A}' \mathbb{J}_m)^{-1} = \mathbb{J}_m \mathbf{A}'^{-1}$. Then, using the relations on $\mathbb{Z}_{m,1}$ given above, we obtain

$$\Delta_{\mathbb{Z}_{m,1}^t, \mathbb{Z}_{m,0}}(\mathbf{A}'^{-1}) = \mathbb{Z}_{m,1}^t \mathbb{J}_m \nabla_{\mathbb{Z}_{m,1}^t, \mathbb{Z}_{m,0}}(\mathbb{J}_m \mathbf{A}'^{-1})$$

and thus a $\Delta_{\mathbb{Z}_{m,1}^t, \mathbb{Z}_{m,0}}$ -generator of $\mathbf{A}'^{-1}$. The general formula above further leads to

$$\begin{aligned} \Delta_{\mathbb{M}_{\mathbf{Q}}^t, \mathbb{M}_{\mathbf{P}}}(\mathbf{R} \mathbf{A}'^{-1} \mathbf{L}) &= -\nabla_{\mathbb{M}_{\mathbf{Q}}^t, \mathbb{Z}_{m,1}^t}(\mathbf{R}) \mathbf{A}'^{-1} \mathbb{Z}_{m,0} \mathbf{L} + \mathbf{R} \Delta_{\mathbb{Z}_{m,1}^t, \mathbb{Z}_{m,0}}(\mathbf{A}'^{-1}) \mathbf{L} \\ &\quad + \mathbb{M}_{\mathbf{Q}}^t \mathbf{R} \mathbf{A}'^{-1} \nabla_{\mathbb{Z}_{m,0}, \mathbb{M}_{\mathbf{P}}}(\mathbf{L}); \end{aligned}$$

since $\mathbf{R} \mathbf{A}'^{-1} \mathbf{L} = \mathbf{A}^{-1}$, this gives a $\Delta_{\mathbb{M}_{\mathbf{Q}}^t, \mathbb{M}_{\mathbf{P}}}$ -generator of $\mathbf{A}^{-1}$ whose length $\alpha + 2$ is finally reduced to α .

In all cases Lemma 24 can be reused, and the requested complexity bounds are then derived in the same way as for the three cases above, up to a minor difference: we also need to take into account the cost of multiplication by $\mathbb{M}_{\mathbf{Q}}$ or its transpose with a vector. However, due to the sparse nature of this matrix, this cost is easily seen to be $O(n)$, so it does not impact the asymptotic estimate.

5. Multiplication algorithms. In this section, we prove bounds on the cost of the product $\mathbf{A}\mathbf{B}$, where $\mathbf{A}$ is a structured matrix in $\mathbb{F}^{m \times n}$ and $\mathbf{B}$ is an arbitrary matrix in $\mathbb{F}^{n \times \beta}$; this will prove the claims in Theorems 1 and 2 regarding multiplication.

THEOREM 27. *For any invertible operator $\mathcal{L}$ associated with $(\mathbf{P}, \mathbf{Q})$, we can take*

$$\mathcal{C}_{\text{mul}}(\mathcal{L}, \alpha, \beta) = O\left(\frac{\beta'}{\alpha'} \mathcal{M}'_{\text{mat}}\left(\frac{p}{\alpha'}, \alpha'\right) + \mathcal{D}(\mathbf{P}, \mathbf{Q}) + (\alpha + \beta)\mathcal{C}(\mathbf{P}, \mathbf{Q})\right),$$

with $p = \max(m, n)$, $\alpha' = \min(\alpha, \beta)$, and $\beta' = \max(\alpha, \beta)$.

Using Lemma 22, it will be sufficient to consider multiplication for the basic operators associated with $(\mathbf{P}, \mathbf{Q})$. Thus, throughout this section we let $\mathcal{L} : \mathbb{F}^{m \times n} \rightarrow \mathbb{F}^{m \times n}$ be one of the two operators $\nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ and $\Delta_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$ and, given $(\mathbf{G}, \mathbf{H}) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha}$ and $\mathbf{B} \in \mathbb{F}^{n \times \beta}$, we show how to compute the product $\mathbf{A}\mathbf{B} \in \mathbb{F}^{m \times \beta}$ with $\mathbf{A}$ the unique matrix in $\mathbb{F}^{m \times n}$ such that $\mathcal{L}(\mathbf{A}) = \mathbf{G}\mathbf{H}^t$.

In Subsection 5.1 we deal with the cases covered by Theorem 13; a key step of the proof (corresponding to the case $Q = x^n$) is deferred to Subsection 5.2 due to its length. We then extend our results to all remaining cases in Subsection 5.3.

In view of Theorem 21, it would actually be sufficient to assume that $\mathcal{L}$ is the Hankel operator $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$. This would allow us to bypass most of the derivation in Subsection 5.3; however, such a restriction does not seem to lead to a better cost estimate than our direct approach, which we include for completeness.

5.1. Applying Theorem 13. Let us first assume that $\mathcal{L} = \nabla_{\mathbb{M}_{\mathbf{P}}, \mathbb{M}_{\mathbf{Q}}^t}$. Given $\mathbf{P}$, $\mathbf{Q}$ and the generator $(\mathbf{G}, \mathbf{H})$, and a matrix $\mathbf{B} \in \mathbb{F}^{n \times \beta}$, Theorem 13 shows that $\mathbf{A}\mathbf{B}$ can be computed as follows:

1. compute the polynomials $P, Q, \gamma_k, \eta_k, Q^{-1} \bmod P_i$ for $k \leq \alpha$ and $i \leq d$;
2. compute the $n \times \beta$ matrix $\mathbf{B}' = \mathbb{X}_{\mathbf{Q}} \mathbb{Y}_{\mathbf{Q}} \mathbf{B}$;
3. compute the $m \times \beta$ matrix $\mathbf{C}$ given by

$$\mathbf{C} = \sum_{k \leq \alpha} \mathbb{M}_{\gamma_k, P, n} \mathbb{M}_{\eta_k, Q} \mathbf{B}';$$

4. return the $m \times \beta$ matrix $\mathbb{V}_{\mathbf{P}, \mathbf{Q}} \mathbb{W}_{\mathbf{P}} \mathbf{C}$.

Proceeding as in the proof of Corollary 14, we see that Steps 1, 2, 4 can be completed in time $O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + (\alpha + \beta)\mathcal{C}(\mathbf{P}, \mathbf{Q}))$. We are thus left with analyzing Step 3.

Let $\mathbf{b}_1, \dots, \mathbf{b}_\beta$ be the columns of $\mathbf{B}'$, and for $i = 1, \dots, \beta$ let $B_i = \text{pol}(\mathbf{b}_i) \in \mathbb{F}[x]_n$ be the polynomial whose coefficients are the entries of $\mathbf{b}_i$. In polynomial terms, given $B_1, \dots, B_\beta$ in $\mathbb{F}[x]_n$, we want to compute $C_1, \dots, C_\beta$ in $\mathbb{F}[x]_m$ such that

$$C_i = \sum_{k \leq \alpha} \gamma_k (\eta_k B_i \bmod Q) \bmod P$$

for $i \leq \beta$; then, the m coefficients of C_i will make up the i th column of the matrix $\mathbf{C}$.

To compute C_i , we can compute the sum first, and then reduce it modulo P . The core problem is thus to compute

$$R_i = \sum_{k \leq \alpha} \gamma_k (\eta_k B_i \bmod Q)$$

in $\mathbb{F}[x]_{m+n-1}$ for all $i \leq \beta$; the cost is given in Theorem 31 of Subsection 5.3 below: writing $\alpha' = \min(\alpha, \beta)$ and $\beta' = \max(\alpha, \beta)$, one can compute $R_1, \dots, R_\beta$ in time

$$(9) \quad O\left(\frac{\beta'}{\alpha'} \mathcal{M}'_{\text{mat}}\left(\frac{p}{\alpha'}, \alpha'\right)\right).$$

Finally, since R_i has degree less than $m + n - 1$ and since P has degree m , computing each $C_i = R_i \bmod P$ takes time $O(\mathcal{M}(m) + \mathcal{M}(n)) \subset O(\mathcal{M}(p))$, so we can deduce $C_1, \dots, C_\beta$ from $R_1, \dots, R_\beta$ for a negligible total cost of $O(\alpha \mathcal{M}(p))$.

In the case $\mathcal{L} = \Delta_{\mathbb{M}_P, \mathbb{M}_Q^t}$, the approach is almost the same. A minor difference is that we replace $\mathbb{V}_{P, Q}$ by $\mathbb{V}'_{P, Q}$: this has no influence on the cost. The other difference is the matrix $\mathbb{J}_n$ appearing in the inversion formula for $\Delta_{\mathbb{M}_P, \mathbb{M}_Q^t}$ (second part of Theorem 13); in polynomial terms, this now leads us to compute the sums

$$R'_i = \sum_{k \leq \alpha} \gamma_k \text{rev}(\eta_k B_i \bmod Q, n - 1)$$

for $i = 1, \dots, \beta$, and then to reduce each of them modulo P .

This problem is actually an instance of the question treated above, for we have

$$\text{rev}(R'_i, m + n - 2) = \sum_{k \leq \alpha} \text{rev}(\gamma_k, m - 1)(\eta_k B_i \bmod Q).$$

Since knowing $\text{rev}(R'_i, m + n - 2)$ gives us R'_i for free, we can appeal here as well to Theorem 31 to compute R'_i , using $\text{rev}(\gamma_k, m - 1)$ instead of γ_k . The cost estimate is the same as before, namely, (9), thus leading to Theorem 27 regarding a product of the form AB for the operator $\mathcal{L} = \Delta_{\mathbb{M}_P, \mathbb{M}_Q^t}$.

5.2. Computing the R_i 's when $Q = x^n$. The previous subsection has shown that for both Sylvester's and Stein's displacement operators, computing the product AB essentially reduces to the following basic problem: given as input

- a monic polynomial Q of degree n ,
- polynomials $U_1, \dots, U_\alpha$ in $\mathbb{F}[x]_m$ and $V_1, \dots, V_\alpha, W_1, \dots, W_\beta$ in $\mathbb{F}[x]_n$,

compute

$$R_i = \sum_{k \leq \alpha} U_k (V_k W_i \bmod Q) \quad \text{for } i = 1, \dots, \beta.$$

The naive algorithm computes all sums independently; writing $p = \max(m, n)$, its cost is $O(\alpha \beta \mathcal{M}(p))$. In this subsection we establish the following improved complexity estimate in the special case where $Q = x^n$. (This result will be further extended to the case of a general Q of degree n in Subsection 5.3.)

THEOREM 28. *Assume $Q = x^n$ and $\alpha \leq n$, and let $p = \max(m, n)$, $\alpha' = \min(\alpha, \beta)$, and $\beta' = \max(\alpha, \beta)$. Then one can compute the polynomials $R_1, \dots, R_\beta$ defined above in time*

$$O\left(\frac{\beta'}{\alpha'} \mathcal{M}'_{\text{mat}}\left(\frac{p}{\alpha'}, \alpha'\right)\right).$$

To prove Theorem 28, we first rephrase our problem in polynomial matrix terms. Let

$$U \in \mathbb{F}[x]_m^{\alpha \times 1}, \quad V \in \mathbb{F}[x]_n^{\alpha \times 1}, \quad W \in \mathbb{F}[x]_n^{\beta \times 1}$$

be the polynomial vectors with respective entries (U_i) , (V_i) and (W_i) . Then, our problem amounts to computing the polynomial row vector $R \in \mathbb{F}[x]_{m+n-1}^{1 \times \beta}$ such that

$$R = U^t (V W^t \bmod x^n).$$

We first solve that problem in the special case where $\alpha = \beta$ (this is where most of the difficulty takes place); then, we reduce the other cases of arbitrary α, β to this case.

Case $\alpha = \beta$. The following lemma is the basis of our algorithm in this case; it is taken from the proof of [4, Lemma 12].

LEMMA 29. Let α, γ, ν be in $\mathbb{N}_{>0}$ with ν even. Let V, W be in $\mathbb{F}[x]_\nu^{\alpha \times \gamma}$ and define

$$V_0 = V \bmod x^{\nu/2}, \quad V_1 = V \operatorname{div} x^{\nu/2},$$

$$W_0 = W \bmod x^{\nu/2}, \quad W_1 = W \operatorname{div} x^{\nu/2}.$$

Then the matrices $[V_0 \ V_1]$ and $[W_1 \ W_0]$ are in $\mathbb{F}[x]_{\nu/2}^{\alpha \times 2\gamma}$, and we have

$$VW^t \bmod x^\nu = V_0 W_0^t + x^{\nu/2} ([V_0 \ V_1] [W_1 \ W_0]^t \bmod x^{\nu/2}).$$

We will want to apply this lemma recursively, starting from $\nu = n$ and $\gamma = 1$. It will be convenient to have both n and α be powers of two.

- If n is not a power of two, we define $\bar{n} = 2^{\lceil \log n \rceil}$ and $\delta = \bar{n} - n$. One may then check that $a \bmod b = x^{-\delta} ((x^\delta a) \bmod (x^\delta b))$ for any a and nonzero b in $\mathbb{F}[x]$; applying this identity componentwise to the definition of R gives

$$(10) \quad R = x^{-\delta} U^t (V(x^\delta W)^t \bmod x^{\bar{n}}),$$

where V and $x^\delta W$ have degree less than $\bar{n}$.

- If α is not a power of two, we define $\bar{\alpha} = 2^{\lceil \log \alpha \rceil}$ and $\mu = \bar{\alpha} - \alpha$. Then, we can introduce μ dummy polynomials (U_i) , (V_i) and (W_i) , all equal to zero, without affecting the value of $R_1, \dots, R_\alpha$.

The resulting algorithm is given in Figure 1, with a top-level procedure `mul` that handles the cases when n or α are not powers of two by defining $\bar{n}$ and $\bar{\alpha}$ as above, and a main recursive procedure `mul_rec`. We initially set $\nu = \bar{n}$ and $\gamma = 1$; through each recursive call, the width γ of the matrices V and W is doubled, while their degree ν is divided by two, so that the invariant $\gamma\nu = \bar{n}$ is maintained.

LEMMA 30. Let $p = \max(m, n)$. Algorithm `mul` of Figure 1 works correctly in time

$$O\left(\mathcal{M}'_{\text{mat}}\left(\frac{p}{\alpha}, \alpha\right)\right).$$

Proof. Correctness is a direct consequence of Lemma 29 and the discussion that followed it. For the cost analysis, we let $C(k)$ denote the cost of `mul_rec` called upon parameters $\bar{\alpha}$ and γ such that $\bar{\alpha} = 2^k \gamma$. Note that the total cost of `mul` will then be at most $C(\kappa)$, with $\kappa = \lceil \log \alpha \rceil$. In particular, below, we always have $k \leq \kappa$. Another useful remark is that we always have $\nu \leq 2p/\gamma$, since $\nu = \bar{n}/\gamma$ and $\bar{n} \leq 2n \leq 2p$.

If $k = 0$, we have $\gamma = \bar{\alpha}$, so we are in the base case of the algorithm, where V and W are in $\mathbb{F}[x]_\nu^{\bar{\alpha} \times \bar{\alpha}}$. We first compute $R' = VW^t \bmod x^\nu$ in time $\mathcal{M}_{\text{mat}}(\nu, \bar{\alpha})$. Then, given U in $\mathbb{F}[x]_m^{\bar{\alpha} \times 1}$ and R' in $\mathbb{F}[x]_\nu^{\bar{\alpha} \times \bar{\alpha}}$, we obtain $R = U^t R'$ in two steps as follows: rewriting U^t in the form $U^t = [1 \ x^c \ x^{2c} \ \dots \ x^{(\bar{\alpha}-1)c}] U'$ with $c = \lfloor m/\bar{\alpha} \rfloor$ and $U' \in \mathbb{F}[x]_c^{\bar{\alpha} \times \bar{\alpha}}$, we compute first $Q = U' R'$ in time $\mathcal{M}_{\text{mat}}(\max\{c, \nu\}, \bar{\alpha})$, and then deduce R from Q using at most $\bar{\alpha}^2(c + \nu - 1)$ additions in $\mathbb{F}$. In summary,

$$\begin{aligned} C(0) &\leq \mathcal{M}_{\text{mat}}(\nu, \bar{\alpha}) + \mathcal{M}_{\text{mat}}(\max\{c, \nu\}, \bar{\alpha}) + \bar{\alpha}^2(c + \nu - 1) \\ &\leq 2\mathcal{M}_{\text{mat}}\left(\frac{2p}{\bar{\alpha}}, \bar{\alpha}\right) + 3\bar{\alpha}p, \end{aligned}$$

using $c \leq m/\bar{\alpha} \leq p/\bar{\alpha}$ and $\nu = \bar{n}/\bar{\alpha} \leq 2p/\bar{\alpha}$. Therefore, there is a constant c_0 with

$$(11) \quad C(0) \leq c_0 \cdot \mathcal{M}_{\text{mat}}\left(\frac{p}{\bar{\alpha}}, \bar{\alpha}\right).$$

```

Algorithm mul_rec(U, V, W, m, ν, ᾱ, γ)
Input: U ∈ ℔[x]_m^{ᾱ×1}, V, W ∈ ℔[x]_ν^{ᾱ×γ}
Assumptions: ᾱ, γ and ν are powers of two and γ ≤ ᾱ
Output: R ∈ ℔[x]_{m+ν-1}^{1×ᾱ} such that R = U^t(VW^t mod x^ν).

if γ = ᾱ then
    R' := VW^t mod x^ν
    R := U^t R'
else
    ν' := ν/2; γ' = 2γ
    V_0 := V mod x^{ν'}; V_1 := V div x^{ν'}; V' := [V_0 V_1]
    W_0 := W mod x^{ν'}; W_1 := W div x^{ν'}; W' := [W_1 W_0]
    R' := mul_rec(U, V', W', m, ν', ᾱ, γ')
    R := U^t V_0 W_0^t + x^{ν'} R'
return R.

```

```

Algorithm mul(U, V, W, m, n, α)
Input: U ∈ ℔[x]_m^{α×1}, V, W ∈ ℔[x]_n^{α×1}
Assumption: α ≤ n
Output: R ∈ ℔[x]_{m+n-1}^{1×α} such that R = U^t(VW^t mod x^n).

n̄ := 2^{⌈log(n)⌉}; δ := n̄ - n
ᾱ := 2^{⌈log(α)⌉}; μ := ᾱ - α
Ū := U augmented with μ zero rows
V̄ := V augmented with μ zero rows
W̄ := x^δ W augmented with μ zero rows
R̄ := mul_rec(Ū, V̄, W̄, m, n̄, ᾱ, 1)
R := x^{-δ} R̄
return the first α entries of R.

```

FIG. 1. Algorithms mul and mul_rec.

Let us now bound $C(k)$ when $k \geq 1$. Given $V, W \in \mathbb{F}[x]_{\nu}^{\alpha \times \gamma}$, we can compute V_0, V_1, W_0, W_1 for free. Then R' is computed recursively in time $C(k-1)$, and it remains to bound the cost of producing the result as $R = Q + x^{\nu'} R'$, with $Q = U^t V_0 W_0^t$.

For now let us write $D(k)$ for the cost of computing Q . Given R' in $\mathbb{F}[x]_{m+\nu'-1}^{1 \times \alpha}$ with $\nu' = \nu/2$ and Q in $\mathbb{F}[x]_{m+\nu-1}^{1 \times \alpha}$, we can add them together using at most $\alpha(m + \nu' - 2)$ additions in $\mathbb{F}$. Since $\nu' = \bar{n}/2\gamma \leq n/\gamma \leq n$, we deduce that for $k \geq 1$,

$$(12) \quad C(k) \leq C(k-1) + D(k) + 2\bar{\alpha}p.$$

In order to bound $D(k)$ let us rewrite as before U^t as $U^t = [1 \ x^c \ x^{2c} \ \dots \ x^{(\gamma-1)c}]U'$ with $c = \lfloor m/\gamma \rfloor$ and $U' \in \mathbb{F}[x]_c^{\gamma \times \alpha}$; we compute Q as $Q = [1, x^c, x^{2c}, \dots, x^{(\gamma-1)c}](U'V_0W_0^t)$. The product $U'V_0W_0^t$ involves three polynomial matrices of respective dimensions $\gamma \times \bar{\alpha}, \bar{\alpha} \times \gamma, \gamma \times \bar{\alpha}$, with entries of degree less than $\max\{c, \nu'\}$. Furthermore, since $c \leq m/\gamma$ and, as seen above, $\nu' \leq n/\gamma$, we have $\max\{c, \nu'\} \leq p/\gamma$. Since $\gamma \leq \bar{\alpha}$, we can proceed as in the second case considered in the proof of [4, Lemma 7] to show that $U'V_0W_0^t$ can be evaluated as $(U'V_0)W_0^t$ in time

$$\frac{\bar{\alpha}}{\gamma} \mathcal{M}_{\text{mat}}\left(\frac{p}{\gamma}, \gamma\right) + 2\bar{\alpha}p + \frac{\bar{\alpha}}{\gamma} \mathcal{M}_{\text{mat}}\left(\frac{2p}{\gamma}, \gamma\right).$$

Since $U^t V_0 W_0^t$ has dimensions $\gamma \times \bar{\alpha}$ and degree less than $3p/\gamma$, reconstructing Q from that product requires at most $\gamma \bar{\alpha} \cdot 3p/\gamma = 3\bar{\alpha}p$ additions in $\mathbb{F}$. Thus, using $\bar{\alpha} = 2^k \gamma$,

$$D(k) \leq 2^k \left(\mathcal{M}_{\text{mat}} \left(\frac{2^k p}{\bar{\alpha}}, \frac{\bar{\alpha}}{2^k} \right) + \mathcal{M}_{\text{mat}} \left(\frac{2^{k+1} p}{\bar{\alpha}}, \frac{\bar{\alpha}}{2^k} \right) \right) + 5\bar{\alpha}p.$$

Combining this bound with (12) and the fact that $\mathcal{M}_{\text{mat}}(2d, n) = O(\mathcal{M}_{\text{mat}}(d, n))$, we deduce that there exists a constant c_1 such that for all $k \geq 1$,

$$(13) \quad C(k) \leq C(k-1) + c_1 \cdot 2^k \mathcal{M}_{\text{mat}} \left(\frac{2^k p}{\bar{\alpha}}, \frac{\bar{\alpha}}{2^k} \right).$$

Taking all $k = 0, \dots, \kappa$ into account, we deduce from (11) and (13) that the total time of `mul` is bounded as

$$C(\kappa) \leq \max\{c_0, c_1\} \cdot \sum_{k=0}^{\kappa} 2^k \mathcal{M}_{\text{mat}} \left(\frac{2^k p}{\bar{\alpha}}, \frac{\bar{\alpha}}{2^k} \right) \in O \left(\mathcal{M}'_{\text{mat}} \left(\frac{p}{\bar{\alpha}}, \bar{\alpha} \right) \right).$$

The conclusion follows from the facts that $p/\bar{\alpha} \leq p/\alpha$ and that $\bar{\alpha} < 2\alpha$. $\square$

Case $\alpha < \beta$. In this situation, we split the vector W into vectors $W_1, \dots, W_c$, with each W_i in $\mathbb{F}[x]_n^{\alpha \times 1}$ and $c = \lceil \beta/\alpha \rceil$ (so W_c may be padded with zeros). Then, algorithm `mul` is applied c times (namely, to the (U, V, W_i)) and by Lemma 30 we obtain $R_1, \dots, R_\beta$ in time

$$O \left(\frac{\beta}{\alpha} \mathcal{M}'_{\text{mat}} \left(\frac{p}{\alpha}, \alpha \right) \right).$$

Case $\beta < \alpha$. In this case, we split the vectors U and V into $U_1, \dots, U_c$ and $V_1, \dots, V_c$, with each U_i and V_i in $\mathbb{F}[x]_n^{\beta \times 1}$ and $c = \lceil \alpha/\beta \rceil$. As before, algorithm `mul` is applied c times, but now to the (U_i, V_i, W) , for a cost of

$$O \left(\frac{\alpha}{\beta} \mathcal{M}'_{\text{mat}} \left(\frac{p}{\beta}, \beta \right) \right);$$

the c results thus produced are then added, for a negligible cost of $O(\alpha p)$.

In the next subsection, we denote by `mul`($U, V, W, m, n, \alpha, \beta$) the algorithm that handles all possible cases for $\alpha \leq n$, following the discussion in the above paragraphs.

5.3. Computing the R_i 's in the general case. We now address the case of an arbitrary Q . Given such a Q of degree n in $\mathbb{F}[x]$, as well as $U_1, \dots, U_\alpha$ in $\mathbb{F}[x]_m$, $V_1, \dots, V_\alpha$ in $\mathbb{F}[x]_n$ and $W_1, \dots, W_\beta$ in $\mathbb{F}[x]_n$, we recall that our goal is to compute

$$R_i = \sum_{k \leq \alpha} U_k (V_k W_i \bmod Q), \quad i = 1, \dots, \beta.$$

THEOREM 31. *Assume $\alpha \leq n$, and let $p = \max(m, n)$, $\alpha' = \min(\alpha, \beta)$, and $\beta' = \max(\alpha, \beta)$. Then one can compute the polynomials $R_1, \dots, R_\beta$ in time*

$$O \left(\frac{\beta'}{\alpha'} \mathcal{M}'_{\text{mat}} \left(\frac{p}{\alpha'}, \alpha' \right) \right).$$

To prove Theorem 31, the idea is to use the well-known fact that Euclidean division reduces to power series inversion and multiplication; in this way, the proof of Theorem 31 will reduce to that of Theorem 28.

The first step is to replace remainders by quotients. By applying the Euclidean division equality $V_k W_i = (V_k W_i \operatorname{div} Q)Q + (V_k W_i \operatorname{mod} Q)$ and defining

$$S_i = \sum_{k \leq \alpha} U_k (V_k W_i \operatorname{div} Q) \quad \text{for } i \leq \beta \quad \text{and} \quad T = \sum_{k \leq \alpha} U_k V_k,$$

we easily deduce the following lemma.

LEMMA 32. $R_i = TW_i - QS_i$ for $i \leq \beta$.

The main issue in our algorithm will be the computation of $S_1, \dots, S_\beta$; we will actually compute their reverse polynomials. Observe that U_i has degree at most $m-1$, $V_k W_i \operatorname{div} Q$ has degree at most $n-2$, and S_i has degree at most $m+n-3$. We thus introduce the polynomials $\tilde{S}_i = \operatorname{rev}(S_i, m+n-3)$; knowing those, we can recover the polynomials S_i for free. For $k \leq \alpha$ and $i \leq \beta$, let us also define

$$\widetilde{U}_k = \operatorname{rev}(U_k, m-1), \quad \widetilde{V}_k = \frac{\operatorname{rev}(V_k, n-1)}{\operatorname{rev}(Q, n)} \operatorname{mod} x^{n-1}, \quad \widetilde{W}_i = \operatorname{rev}(W_i, n-1) \operatorname{mod} x^{n-1}.$$

All these polynomials are related by the following formula.

LEMMA 33. $\tilde{S}_i = \sum_{k \leq \alpha} \widetilde{U}_k (\widetilde{V}_k \widetilde{W}_i \operatorname{mod} x^{n-1})$ for $i \leq \beta$.

Proof. Using that $\operatorname{rev}(ab, r+s) = \operatorname{rev}(a, r) \operatorname{rev}(b, s)$ for any polynomials $a, b \in \mathbb{F}[x]$ of respective degrees r and s , we deduce from the definitions of S_i and $\tilde{S}_i$ that

$$\tilde{S}_i = \sum_{k \leq \alpha} \widetilde{U}_k \operatorname{rev}(V_k W_i \operatorname{div} Q, n-2).$$

Following [12, p. 258], we conclude with the equalities

$$\begin{aligned} \operatorname{rev}(V_k W_i \operatorname{div} Q, n-2) &= \frac{\operatorname{rev}(V_k W_i, 2n-2)}{\operatorname{rev}(Q, n)} \operatorname{mod} x^{n-1} \\ &= \frac{\operatorname{rev}(V_k, n-1) \operatorname{rev}(W_i, n-1)}{\operatorname{rev}(Q, n)} \operatorname{mod} x^{n-1} \\ &= \widetilde{V}_k \widetilde{W}_i \operatorname{mod} x^{n-1}. \end{aligned} \quad \square$$

Figure 2 details the algorithm we just sketched. To prove Theorem 31, it suffices to observe that the cost boils down to one call to mul plus $O((\alpha + \beta) \mathcal{M}(p))$ for all other operations, and then to apply Theorem 28.

6. Algorithms for inversion and system solving. We conclude in this section by establishing the cost bounds announced in Theorem 1 for structured inversion and structured system solving. We rely on Theorem 21 from Section 4 to reduce our task to the case of the Hankel operator $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$, then use the fast multiplication algorithm of Section 5 to speed up the Morf/Bitmead-Anderson (MBA) approach.

THEOREM 34. For any invertible operator $\mathcal{L}$ associated with $(\mathbf{P}, \mathbf{Q})$, we can take

$$\mathcal{C}_{\operatorname{inv}}(\mathcal{L}, \alpha) = O\left(\mathcal{M}_{\operatorname{mat}}''\left(\frac{m}{\alpha}, \alpha\right)\right)$$

and

$$\mathcal{C}_{\operatorname{solve}}(\mathcal{L}, \alpha) = O\left(\mathcal{M}_{\operatorname{mat}}''\left(\frac{p}{\alpha}, \alpha\right)\right), \quad p = \max(m, n).$$

Algorithm $\text{mulQ}(\mathbf{U}, \mathbf{V}, \mathbf{W}, m, n, \alpha, \beta, Q)$
Input: $\mathbf{U} \in \mathbb{F}[x]_m^{\alpha \times 1}$, $\mathbf{V} \in \mathbb{F}[x]_n^{\alpha \times 1}$, $\mathbf{W} \in \mathbb{F}[x]_n^{\beta \times 1}$
 Q monic of degree n in $\mathbb{F}[x]$
Assumption: $\alpha \leq n$
Output: $\mathbf{R} \in \mathbb{F}[x]_{m+n-1}^{1 \times \beta}$ such that $\mathbf{R} = \mathbf{U}^t(\mathbf{V}\mathbf{W}^t \bmod Q)$.
 $\tilde{\mathbf{U}} := [\text{rev}(U_k, m-1), k=1, \dots, \alpha]$
 $\tilde{\mathbf{V}} := [\text{rev}(V_k, n-1)/\text{rev}(Q, n) \bmod x^{n-1}, k=1, \dots, \alpha]$
 $\tilde{\mathbf{W}} := [\text{rev}(W_i, n-1) \bmod x^{n-1}, i=1, \dots, \beta]$
 $\tilde{\mathbf{S}} := \text{mul}(\tilde{\mathbf{U}}, \tilde{\mathbf{V}}, \tilde{\mathbf{W}}, m, n-1, \alpha, \beta)$
 $\mathbf{T} := \sum_{k=1, \dots, \alpha} U_k V_k$
 $\mathbf{R} := [T W_i - Q \text{rev}(\tilde{S}_i, m+n-3), i=1, \dots, \beta]$
 return $\mathbf{R}$.

FIG. 2. Algorithm mulQ .

Recall that Theorem 21 in Section 4 proves that we can take

$$\mathcal{C}_{\text{inv}}(\mathcal{L}, \alpha) \leq \mathcal{C}_{\text{inv}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{m,1}^t}, \alpha + 2) + O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha \mathcal{C}(\mathbf{P}, \mathbf{Q}) + \alpha^{\omega-1} m)$$

and

$$\mathcal{C}_{\text{solve}}(\mathcal{L}, \alpha) \leq \mathcal{C}_{\text{solve}}(\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}, \alpha + 2) + O(\mathcal{D}(\mathbf{P}, \mathbf{Q}) + \alpha \mathcal{C}(\mathbf{P}, \mathbf{Q})),$$

so we are left with estimating the cost of inversion and system solving for $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$. For this, we shall use the MBA approach with preconditioning introduced in [20] together with the “compression-free” formulas from [16] for generating the matrix inverse. Such formulas being expressed with products of the form (structured matrix) $\times$ (unstructured matrix), this provides a way to exploit our results for fast structured matrix multiplication from Section 5. Specifically, starting from a $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$ -generator $(\mathbf{G}, \mathbf{H})$ of length α of $\mathbf{A} \in \mathbb{F}^{m \times n}$, we proceed in three steps as follows.

Preconditioning. From $(\mathbf{G}, \mathbf{H})$ we begin by deducing a generator of the preconditioned matrix

$$\tilde{\mathbf{A}} = \mathbb{U}(\mathbf{v}_1) \mathbf{A} \mathbb{U}(\mathbf{v}_2)^t,$$

where each $\mathbb{U}(\mathbf{v}_i)$ is a unit upper triangular Toeplitz matrix defined by some random vector $\mathbf{v}_i$; as shown in [21], such a preconditioning ensures that, with high probability, $\tilde{\mathbf{A}}$ has *generic rank profile*, that is, denoting by $\tilde{\mathbf{A}}_k$ the leading principal submatrix of $\tilde{\mathbf{A}}$ of dimensions $k \times k$,

$$\det(\tilde{\mathbf{A}}_k) \neq 0 \quad \text{for } k = 1, \dots, \text{rank}(\tilde{\mathbf{A}}).$$

In our case, preconditioning will be done in such a way that the generator of $\tilde{\mathbf{A}}$ corresponds not to $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}$ but to another Hankel operator, namely $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,0}^t}$ (algorithm `precond`, Figure 3). The latter operator being singular but still *partly regular* [33, §4.5], such a generator will in fact be a triple $(\tilde{\mathbf{G}}, \tilde{\mathbf{H}}, \tilde{\mathbf{u}})$, with $(\tilde{\mathbf{G}}, \tilde{\mathbf{H}})$ of length $O(\alpha)$ and $\tilde{\mathbf{u}}$ carrying the last row of $\tilde{\mathbf{A}}$. Also, we will guarantee that the matrices $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{H}}$ have the following special shape: their first α columns are precisely $\mathbb{U}(\mathbf{v}_1)\mathbf{G}$ and $\mathbb{U}(\mathbf{v}_2)\mathbf{H}$. These two requirements will be key to exploit the “compression-free” inversion scheme of [16, §4.3] and then to move back from, say, $\tilde{\mathbf{A}}^{-1}$ to $\mathbf{A}^{-1}$.

Computation of the leading principal inverse. Assuming that $\tilde{A}$ has generic rank profile and given $(\tilde{G}, \tilde{H}, \tilde{u})$, we now consider computing a compact representation of the inverse of its largest nonsingular leading principal submatrix, that is, of $\tilde{A}_r^{-1}$ with

$$r = \text{rank}(\tilde{A}) = \text{rank}(A).$$

Since $\tilde{A}_r$ is structured with respect to $\nabla_{\mathbb{Z}_{r,0}, \mathbb{Z}_{r,0}^t}$, its inverse is structured with respect to $\nabla_{\mathbb{Z}_{r,0}^t, \mathbb{Z}_{r,0}}$ and, as recalled in (1),

$$\text{rank}(\nabla_{\mathbb{Z}_{r,0}^t, \mathbb{Z}_{r,0}}(\tilde{A}_r^{-1})) = \text{rank}(\nabla_{\mathbb{Z}_{r,0}, \mathbb{Z}_{r,0}^t}(\tilde{A})) =: \rho.$$

Therefore, our second step (detailed in Section 6.2) will be to compute a $\nabla_{\mathbb{Z}_{r,0}^t, \mathbb{Z}_{r,0}}$ -generator of $\tilde{A}_r^{-1}$ of length ρ .

In fact, following [33, §5], we give an algorithm (`lp_inv`, Figure (6)) that does *not* assume $\tilde{A}$ has generic rank profile but discovers whether this is the case or not. It calls first an auxiliary routine, called `largest`, which returns the size ℓ of the largest leading principal submatrix $\tilde{A}_\ell$ having generic rank profile together with a generator of the inverse $\tilde{A}_\ell^{-1}$. Then, from such a generator, one can check efficiently whether ℓ equals $\text{rank}(\tilde{A})$, which is a condition equivalent to $\tilde{A}$ having generic rank profile. Thus, overall, algorithm `lp_inv` generates $\tilde{A}_r^{-1}$ if and only if $\tilde{A}$ has generic rank profile (and reports 'failure' otherwise, since this provides a way of certifying whether preconditioning A was successful or not).

Algorithm `largest` calls a core recursive routine `largest_rec`, which can be seen as a combination of Kaltofen's algorithm Leading Principal Inverse [20] and [16, algorithm `GenInvHL`], which thus relies on products of the form (structured matrix) $\times$ (unstructured matrix). Also, the generating matrices $\tilde{Y}_\ell$ and $\tilde{Z}_\ell$ produced by algorithm `largest` are specified to have the following shape:

$$\tilde{Y}_\ell = -\tilde{A}_\ell^{-1}\tilde{G}_\ell, \quad \tilde{Z}_\ell = \tilde{A}_\ell^{-t}\tilde{H}_\ell,$$

where $\tilde{G}_\ell$ is made from the first ℓ rows of $\tilde{G}$, and similarly for $\tilde{H}_\ell$. When $\ell = r$, the same specification is inherited by the generator $(\tilde{Y}_r, \tilde{Z}_r, \tilde{v})$ produced by algorithm `lp_inv`, whose matrices satisfy

$$\tilde{Y}_r = -\tilde{A}_r^{-1}\tilde{G}_r, \quad \tilde{Z}_r = \tilde{A}_r^{-t}\tilde{H}_r$$

where $\tilde{G}_r$ and $\tilde{H}_r$ are the first r rows of $\tilde{G}$ and $\tilde{H}$, and $\tilde{v}^t$ is the first row of $\tilde{A}_r^{-1}$.

Generating inverses and solving linear systems. Given the rank r of $\tilde{A}$ and the $\nabla_{\mathbb{Z}_{r,0}^t, \mathbb{Z}_{r,0}}$ -generator $(\tilde{Y}_r, \tilde{Z}_r, \tilde{v})$ of $\tilde{A}_r^{-1}$ as above, it is immediate to decide whether $\tilde{A}$ and A are invertible and, if so, to deduce the $\nabla_{\mathbb{Z}_{m,0}^t, \mathbb{Z}_{m,1}}$ -generator of A^{-1} given by

$$Y = -A^{-1}G, \quad Z = A^{-t}H.$$

This corresponds to algorithm `inv` in Figure 7. Now, given an additional vector $\mathbf{b} \in \mathbb{F}^m$, we reduce the study of $Ax = \mathbf{b}$ to that of the equivalent linear system $\tilde{A}\tilde{x} = \tilde{\mathbf{b}}$, where $\tilde{\mathbf{b}} = \mathbb{U}(\mathbf{v}_1)\mathbf{b}$ and for which an algorithm can be derived directly from the one in [21, §4], with the additional guarantee that if $\mathbf{b} = 0$ and the column rank of A is not full, then a nonzero solution $\mathbf{x}$ is obtained. This corresponds to algorithm `solve` in

Figure 8. (Clearly, both `inv` and `solve` are Las Vegas algorithms—“always correct, probably fast”, thanks to the specification of algorithm `lp_inv`.)

The next three sections provide detailed descriptions of the algorithms mentioned above, namely `precond`, `largest_rec`, `largest`, `lp_inv`, `inv`, `solve`, together with their correctness and complexity proofs, thereby establishing Theorem 34.

6.1. Preconditioning. We precondition our structured matrix A as shown in Figure 3 below. Here and hereafter, $\mathbf{e}_{n,i}$ denotes the i th unit vector in $\mathbb{K}^n$ and, for any given $m \times n$ matrix M such that $n \geq \alpha$, we write $M^{\rightarrow \alpha}$ for the $m \times \alpha$ matrix obtained by keeping only the first α columns of M .

Algorithm `precond`($G, H, \mathbf{v}_1, \mathbf{v}_2$)

Input: $(G, H) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha}$ such that $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}(A) = GH^t$,
 $\mathbf{v}_1 \in \mathbb{F}^m$ and $\mathbf{v}_2 \in \mathbb{F}^n$.

Output: $(\tilde{G}, \tilde{H}) \in \mathbb{F}^{m \times (\alpha+4)} \times \mathbb{F}^{n \times (\alpha+4)}$ and $\tilde{\mathbf{u}} \in \mathbb{F}^n$ such that

- $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,0}^t}(\tilde{A}) = \tilde{G}\tilde{H}^t$ for $\tilde{A} = \mathbb{U}(\mathbf{v}_1)A\mathbb{U}(\mathbf{v}_2)^t$;
- $\tilde{\mathbf{u}}^t$ is the last row of $\tilde{A}$;
- $\tilde{G}^{\rightarrow \alpha} = \mathbb{U}(\mathbf{v}_1)G$ and $\tilde{H}^{\rightarrow \alpha} = \mathbb{U}(\mathbf{v}_2)H$.

$G_1 := [\mathbb{Z}_{m,0}\mathbb{J}_m\mathbf{v}_1 \mid -\mathbf{e}_{m,1}]$; $H_1 := [\mathbf{e}_{m,m} \mid \mathbb{Z}_{m,0}^t\mathbf{v}_1]$
 $G_2 := [\mathbb{Z}_{n,1}^t\mathbf{v}_2 \mid -\mathbf{e}_{n,n}]$; $H_2 := [\mathbf{e}_{n,1} \mid \mathbb{Z}_{n,0}\mathbb{J}_n\mathbf{v}_2]$
 $\tilde{G} := [\mathbb{U}(\mathbf{v}_1)G \mid G_1 \mid \mathbb{U}(\mathbf{v}_1)AG_2]$; $\tilde{H} := [\mathbb{U}(\mathbf{v}_2)H \mid \mathbb{U}(\mathbf{v}_2)A^tH_1 \mid H_2]$
 $\tilde{\mathbf{u}} := \mathbb{U}(\mathbf{v}_2)A^t\mathbb{U}(\mathbf{v}_1)^t\mathbf{e}_{m,m}$
 return $(\tilde{G}, \tilde{H}, \tilde{\mathbf{u}})$.

FIG. 3. Algorithm `precond`.

LEMMA 35. *Algorithm `precond` works correctly in time $O(\alpha M(p))$. Furthermore, if the vectors $\mathbf{v}_1 \in \mathbb{F}^m$ and $\mathbf{v}_2 \in \mathbb{F}^n$ have their first entry equal to 1 and their remaining $m+n-2$ entries chosen uniformly at random from a finite subset $S \subset \mathbb{F}$, then the matrix $\tilde{A} = \mathbb{U}(\mathbf{v}_1)A\mathbb{U}(\mathbf{v}_2)^t$ has generic rank profile with probability at least*

$$1 - r(r+1)/|S|,$$

where $r = \text{rank}(\tilde{A}) = \text{rank}(A)$ and $|S|$ is the cardinality of S .

Proof. We start by checking $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,0}^t}(\tilde{A}) = \tilde{G}\tilde{H}^t$, from which correctness follows. Writing $\mathcal{L}_1 = \nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{m,0}}(\mathbb{U}(\mathbf{v}_1))$ and $\mathcal{L}_2 = \nabla_{\mathbb{Z}_{n,1}^t, \mathbb{Z}_{n,0}^t}(\mathbb{U}(\mathbf{v}_2)^t)$ and applying (8) gives

$$\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,0}^t}(\tilde{A}) = \mathcal{L}_1 A \mathbb{U}(\mathbf{v}_2)^t + \mathbb{U}(\mathbf{v}_1) G H^t \mathbb{U}(\mathbf{v}_2)^t + \mathbb{U}(\mathbf{v}_1) A \mathcal{L}_2,$$

and it remains to check that $\mathcal{L}_1 = G_1 H_1^t$ and $\mathcal{L}_2 = G_2 H_2^t$. Since $\mathbb{U}(\mathbf{v}_1)$ is upper triangular Toeplitz, the matrix $\mathcal{L}_1 = \mathbb{Z}_{m,0} \mathbb{U}(\mathbf{v}_1) - \mathbb{U}(\mathbf{v}_1) \mathbb{Z}_{m,0}$ is zero everywhere except on its last column, which is equal to $\mathbb{Z}_{m,0} \mathbb{J}_m \mathbf{v}_1$, and on its first row, which is equal to $-\mathbf{v}_1^t \mathbb{Z}_{m,0}$. Hence $\mathcal{L}_1 = \mathbb{Z}_{m,0} \mathbb{J}_m \mathbf{v}_1 \mathbf{e}_{m,m}^t - \mathbf{e}_{m,1} (\mathbb{Z}_{m,0}^t \mathbf{v}_1)^t = G_1 H_1^t$. For $\mathcal{L}_2$, we proceed similarly, by noting that $\mathcal{L}_2 = \mathbb{Z}_{n,1}^t \mathbb{U}(\mathbf{v}_2)^t - \mathbb{U}(\mathbf{v}_2)^t \mathbb{Z}_{n,0}^t$ is zero everywhere but on its first column and last row, equal to $\mathbb{Z}_{n,1}^t \mathbf{v}_2$ and $-\mathbb{Z}_{n,0} \mathbb{J}_n \mathbf{v}_2$, respectively.

Let us now bound the cost of deducing $\tilde{G}$, $\tilde{H}$, $\tilde{u}$ from G , H , v_1 , v_2 . First, we set up the $m \times 2$ matrices G_1 and H_1 and the $n \times 2$ matrices G_2 and H_2 in time $O(p)$. Then, since A satisfies $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}(A) = GH^t$, multiplying A or A^t by a single vector can be done in time $O(\alpha M(p))$, using for example the reconstruction formula in [16, (21b)]. Hence we obtain the products AG_2 , $A^t H_1$, and $A^t \mathbb{U}(v_1)^t e_{m,m}$ in time $O(\alpha M(p))$. Finally, since G and H have α columns each, it remains to multiply $O(\alpha)$ vectors by the triangular Toeplitz matrices $\mathbb{U}(v_1)$ and $\mathbb{U}(v_2)$, and this can be done in time $O(\alpha M(p))$. Overall, the cost of the algorithm is thus bounded by $O(\alpha M(p))$.

The probability analysis is due to Kaltofen and Saunders [21, Theorem 2]. $\square$

6.2. Computation of the leading principal inverse. In order to generate the inverse of the largest nonsingular leading principal submatrix of A , we start with the generation of the largest leading principal submatrix having generic rank profile. This is done by algorithms `largest_rec` and `largest` displayed in Figures 4 and 5. In algorithm `largest_rec` the matrices A_{ij} , G_i , H_j for $1 \leq i, j \leq 2$ have dimensions $m_i \times m_j$, $m_i \times \alpha$, $m_j \times \alpha$, respectively, and correspond to the block partitions

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}, \quad G = \begin{bmatrix} G_1 \\ G_2 \end{bmatrix}, \quad H = \begin{bmatrix} H_1 \\ H_2 \end{bmatrix}.$$

Similarly, u_{ij}^t denotes the last row of A_{ij} .

LEMMA 36. *Algorithm `largest_rec` is correct and if $m = n$ is an integer power of two, then its cost is*

$$O\left(\mathcal{M}_{\text{mat}}''\left(\frac{m}{\alpha}, \alpha\right)\right).$$

Proof. Correctness follows directly from combining the analysis in [20, pp. 801–803] with the formulas for G_5 , H_5 , u_5 , Y , Z , w given in [16, p. 287].

Assuming $m = n$ is a power of two, let $C(m, \alpha)$ denote the cost of algorithm `largest_rec`. We begin by showing that we can take

$$C(m, \alpha) = O(\alpha^\omega) \quad \text{if } \alpha \leq m < 2\alpha.$$

Given G and H we compute the product GH^t in time $O(\alpha^\omega)$. Then, starting from the last row of A and using the fact that for $i, j > 1$ the (i, j) entry of GH^t equals $a_{i-1,j} - a_{i,j-1}$, we deduce all the entries of A in time $O(\alpha^2)$.

Let us now bound the cost when $m \geq 2\alpha$. Proceeding as in the proof of [16, Lemma 5], it is easily seen that one can compute some generators of length at most $\alpha + 2$ for the matrices A_{21} , A_{12}^t , $A_{11}^{-1}A_{12}$, $A_{22}^{-t}A_{21}^t$ as well as the vectors u_{11} , u_{21} , u_{22} , u_5 for a total time in $O(\alpha M(m))$. The number of additions is in $O(\alpha m)$ and we also have to perform the four matrix products $A_{21}Y_{11}$, $A_{12}^t Z_{11}$, $(A_{11}^{-1}A_{12})Y_5$, and $(A_{11}^{-T}A_{21}^T)Z_5$. For example, A_{12} has displacement rank at most $\alpha + 2$ with respect to the operator $\nabla_{\mathbb{Z}_{m/2,0}, \mathbb{Z}_{m/2,1}^t}$, so that using a decomposition of the form $A_{12} = A'_{12} + A''_{12}$ with A'_{12} and A''_{12} of displacement ranks at most α and 2, respectively, we can evaluate the product $A_{21}Y_{11}$ in time $\mathcal{C}_{\text{mul}}(\nabla_{\mathbb{Z}_{m/2,0}, \mathbb{Z}_{m/2,1}^t}, \alpha, \alpha) + O(\alpha M(m))$; by Theorem 2, this is in $O\left(\frac{m}{2\alpha} \mathcal{M}'_{\text{mat}}\left(\frac{m}{2\alpha}, \alpha\right)\right)$. Adding up all these costs thus leads to

$$C(m, \alpha) = 2C\left(\frac{m}{2}, \alpha\right) + O\left(\mathcal{M}'_{\text{mat}}\left(\frac{m}{2\alpha}, \alpha\right)\right) \quad \text{if } m \geq 2\alpha.$$

Hence, for some constant c_0 ,

$$C(m, \alpha) \leq c_0 \left(\sum_{i=0}^{i_0-1} 2^i \mathcal{M}'_{\text{mat}}\left(\frac{m}{2^{i+1}\alpha}, \alpha\right) + 2^{i_0} \alpha^\omega \right)$$

Algorithm `largest_rec`(G, H, u)

Input: $(G, H, u) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha} \times \mathbb{F}^n$ such that $\alpha \leq \min(m, n)$ and

$\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,0}^t}(A) = GH^t$ and $u^t = e_{m,m}^t A$ (the last row of A).

Output: $(\ell, Y, Z, v) \in \mathbb{N} \times \mathbb{F}^{\ell \times \alpha} \times \mathbb{F}^{\ell \times \alpha} \times \mathbb{F}^\ell$ such that ℓ is the order of the largest leading principal submatrix A_ℓ of A having generic rank profile, $Y = -A_\ell^{-1} G_\ell$, $Z = A_\ell^{-t} H_\ell$ and $v = A_\ell^{-t} e_{\ell,1}$ (the first row of A_ℓ^{-1}).

if $\min(m, n) < 2\alpha$ then

 compute A explicitly and then deduce ℓ, Y, Z, v .

else

$m_1 := \lceil m/2 \rceil; m_2 := \lfloor m/2 \rfloor; n_1 := \lceil n/2 \rceil; n_2 := \lfloor n/2 \rfloor$

$(\ell_{11}, Y_{11}, Z_{11}, v_{11}) := \text{largest_rec}(G_1, H_1, u_{11})$

 if $\ell_{11} < \min(m_1, n_1)$ then

$(\ell, Y, Z, v) := (\ell_{11}, Y_{11}, Z_{11}, v_{11})$

 else

$G_S := G_2 + A_{21} Y_{11}; H_S := H_2 - A_{12}^t Z_{11}$

$u_S := u_{22} - A_{12}^t A_{11}^{-t} u_{21}$

 if the $(1, 1)$ element of S is zero then

$(\ell, Y, Z, v) := (\ell_{11}, Y_{11}, Z_{11}, v_{11})$

 else

$(\ell_S, Y_S, Z_S, v_S) := \text{largest_rec}(G_S, H_S, u_S)$

$\ell := \ell_{11} + \ell_S$

$Y := \begin{bmatrix} Y_{11} - (A_{11}^{-1} A_{12}) Y_S \\ Y_S \end{bmatrix}; Z := \begin{bmatrix} Z_{11} - (A_{11}^{-T} A_{21}^T) Z_S \\ Z_S \end{bmatrix};$

$w := -S^{-T} A_{12}^T v_{11}; v := \begin{bmatrix} v_{11} - A_{11}^{-T} A_{21}^T w \\ w \end{bmatrix};$

return (ℓ, Y, Z, v) .

FIG. 4. Algorithm `largest_rec`. Here $A_\ell \in \mathbb{F}^{\ell \times \ell}$ denotes the largest leading principal submatrix of A whose rank profile is generic, and G_ℓ and H_ℓ are the $\ell \times \alpha$ submatrices consisting of the first ℓ rows of G and H , respectively.

with $i_0 \in \mathbb{N}$ such that $\alpha \leq m/2^{i_0} < 2\alpha$, that is, $i_0 = \lfloor \log(m/\alpha) \rfloor$. Defining $\bar{\alpha} = 2^{\lceil \log(\alpha) \rceil}$, we have $\alpha \in (\bar{\alpha}/2, \bar{\alpha}]$, which implies $m/\alpha \in [m/\bar{\alpha}, 2m/\bar{\alpha})$ and thus, since m is an integer power of two,

$$2^{i_0} = m/\bar{\alpha}.$$

Besides, $m/(2^{i+1}\alpha) < m/(2^i\bar{\alpha})$, which by assumption on $\mathcal{M}'_{\text{mat}}$ implies

$$\mathcal{M}'_{\text{mat}}\left(\frac{m}{2^{i+1}\alpha}, \alpha\right) = O\left(\mathcal{M}'_{\text{mat}}\left(\frac{m}{2^i\bar{\alpha}}, \alpha\right)\right).$$

Third, $\alpha^\omega = O(\mathcal{M}'_{\text{mat}}(1, \alpha))$. Consequently, for some constant c_1 ,

$$C(m, \alpha) \leq c_1 \cdot \sum_{i=0}^{\log(m/\bar{\alpha})} 2^i \mathcal{M}'_{\text{mat}}\left(\frac{m}{2^i\bar{\alpha}}, \alpha\right).$$

Now, $\alpha \leq \bar{\alpha}$ implies $m/\bar{\alpha} \leq m/\alpha \leq \overline{m/\alpha}$, where $\overline{m/\alpha}$ denotes the smallest integer power of two greater than or equal to m/α . Hence $\log(m/\bar{\alpha}) \leq \log(\overline{m/\alpha})$ and $\mathcal{M}'_{\text{mat}}(2^{-i}m/\bar{\alpha}) = O(\mathcal{M}'_{\text{mat}}(2^{-i}\overline{m/\alpha}))$ by assumption on $\mathcal{M}'_{\text{mat}}$, so that the sum in the above bound on $C(m, \alpha)$ is in $O(\mathcal{M}''_{\text{mat}}(m/\alpha, \alpha))$, as announced. $\square$

Algorithm **largest**(G, H, u)

Input: $(G, H, u) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha} \times \mathbb{F}^n$ such that $\alpha \leq \min(m, n)$ and
 $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,0}^t}(A) = GH^t$ and $u^t = e_{m,m}^t A$ (the last row of A).

Output: $(\ell, Y, Z, v) \in \mathbb{N} \times \mathbb{F}^{\ell \times \alpha} \times \mathbb{F}^{\ell \times \alpha} \times \mathbb{F}^\ell$ such that ℓ is the order of the largest
 leading principal submatrix A_ℓ of A having generic rank profile,
 $Y = -A_\ell^{-1} G_\ell$, $Z = A_\ell^{-t} H_\ell$ and $v = A_\ell^{-t} e_{\ell,1}$ (the first row of A_ℓ^{-1}).

$\bar{p} := 2^{\lceil \log_2(\max(m,n)) \rceil}$

compute $\bar{\alpha} \in \mathbb{N}$, $\bar{G}, \bar{H} \in \mathbb{F}^{\bar{p} \times \bar{\alpha}}$, and $\bar{u} \in \mathbb{F}^{\bar{p}}$ such that:

- $(\bar{G}, \bar{H}, \bar{u})$ is a $\nabla_{\mathbb{Z}_{\bar{p},0}, \mathbb{Z}_{\bar{p},0}^t}$ -generator of length $\bar{\alpha}$ of $\bar{A} = \begin{bmatrix} A & 0 \\ 0 & 0 \end{bmatrix} \in \mathbb{F}^{\bar{p} \times \bar{p}}$;
- $\alpha \leq \bar{\alpha} \leq \min(\alpha + 2, \bar{p})$;
- $\bar{G} = \begin{bmatrix} G & * \\ * & * \end{bmatrix}$ and $\bar{H} = \begin{bmatrix} H & * \\ * & * \end{bmatrix}$
- $\bar{u}^t$ is the last row of $\bar{A}$

$(\ell, \bar{Y}, \bar{Z}, v) := \text{largest_rec}(\bar{G}, \bar{H}, \bar{u})$

$Y := \bar{Y}^{\mapsto \alpha}$; $Z := \bar{Z}^{\mapsto \alpha}$

return (ℓ, Y, Z, v) .

FIG. 5. Algorithm **largest**. Here $A_\ell \in \mathbb{F}^{\ell \times \ell}$ denotes the largest leading principal submatrix of A whose rank profile is generic, and G_ℓ and H_ℓ are the $\ell \times \alpha$ submatrices consisting of the first ℓ rows of G and H , respectively.

LEMMA 37. Let $p = \max(m, n)$. Algorithm **largest** works correctly in time

$$O\left(\mathcal{M}_{\text{mat}}''\left(\frac{p}{\alpha}, \alpha\right)\right).$$

Proof. Let us show first how to generate the augmented matrix $\bar{A}$. If $\bar{p} = m = n$, then $\bar{A} = A$ and it suffices to take $\bar{\alpha} = \alpha$, $\bar{G} = G$, and $\bar{H} = H$. If $\bar{p} = n > m$ then $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,0}^t}(A) = GH^t$ leads to the following generator of length $\bar{\alpha} := \alpha + 1$:

$$\nabla_{\mathbb{Z}_{\bar{p},0}, \mathbb{Z}_{\bar{p},0}^t}(\bar{A}) = \bar{G}\bar{H}^t, \quad \bar{G} = \begin{bmatrix} G & \\ & e_{\bar{p}-m,1} \end{bmatrix}, \quad \bar{H} = [H \mid u];$$

the range constraint on $\bar{\alpha}$ is satisfied, since $\alpha \leq \min(m, n) = m < n = \bar{p}$, and, on the other hand, the upper left corners of $\bar{G}$ and $\bar{H}$ are G and H , respectively. Proceeding similarly when $\bar{p} = m > n$, we take

$$\bar{\alpha} = \alpha + 1, \quad \bar{G} = [G \mid -u'], \quad \bar{H} = \begin{bmatrix} H & \\ & e_{\bar{p}-n,1} \end{bmatrix},$$

where u' is the last column of A and can be obtained in time $O(\alpha M(\bar{p}))$. It remains to handle the case $\bar{p} > \max(m, n)$, where A is bordered by both some zero rows and some zero columns. In this case, we deduce from $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,0}^t}(A) = GH^t$ that a $\nabla_{\mathbb{Z}_{\bar{p},0}, \mathbb{Z}_{\bar{p},0}^t}$ -generator of length $\alpha + 2$ of $\bar{A}$ is given by

$$\bar{G} = \begin{bmatrix} G & & -u' \\ & e_{\bar{p}-m,1} & \end{bmatrix}, \quad \bar{H} = \begin{bmatrix} H & u & \\ & & e_{\bar{p}-n,1} \end{bmatrix}.$$

Again, it is clear that the upper left corners of $\bar{G}$ and $\bar{H}$ are G and H , respectively. Furthermore, if $\bar{p} \geq \alpha + 2$, we can take $\bar{\alpha} = \alpha + 2$; otherwise, since $\bar{p} > \max(m, n) \geq$

$\min(m, n) \geq \alpha$, we are in the situation where $\bar{p} = \alpha + 1$ and $m = n = \alpha$. In this case we shall proceed as follows to reduce the generator length from $\alpha + 2$ to $\alpha + 1$ while ensuring that $\mathbf{G}$ and $\mathbf{H}$ are in the upper left corners of the new generator matrices: the matrices $\bar{\mathbf{G}}$ and $\bar{\mathbf{H}}$ defined above have dimensions $(\alpha + 1) \times (\alpha + 2)$ and are given by

$$\bar{\mathbf{G}} = \begin{bmatrix} \mathbf{G} & & -\mathbf{u}' \\ & 1 & \\ & & \end{bmatrix} \quad \text{and} \quad \bar{\mathbf{H}}^t = \begin{bmatrix} \mathbf{H}^t & \\ \mathbf{u}^t & \\ & 1 \end{bmatrix};$$

since $\mathbf{G}$ is $\alpha \times \alpha$, we can deduce from $\mathbf{G}$ and $\mathbf{u}'$ a vector $\mathbf{u}'' \in \mathbb{F}^\alpha$ such that $\mathbf{G}\mathbf{u}'' = \mathbf{u}'$ in time $O(\alpha^\omega)$; then, using the fact that

$$\mathbf{E} = \begin{bmatrix} \mathbb{I}_\alpha & & \mathbf{u}'' \\ & 1 & \\ & & 1 \end{bmatrix} \implies \bar{\mathbf{G}}\mathbf{E} = \begin{bmatrix} \mathbf{G} & & 0 \\ & 1 & 0 \\ & & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{E}^{-1}\bar{\mathbf{H}}^t = \begin{bmatrix} \mathbf{H}^t & -\mathbf{u}'' \\ \mathbf{u}^t & \\ & 1 \end{bmatrix},$$

we conclude that a suitable $\nabla_{\mathbb{Z}_{\bar{p},0}, \mathbb{Z}_{\bar{p},0}^t}$ -generator of $\bar{\mathbf{A}}$ is obtained by taking the first $\alpha + 1 =: \bar{\alpha}$ columns of $\bar{\mathbf{G}}\mathbf{E}$ and $\bar{\mathbf{H}}\mathbf{E}^{-t}$. Finally, obtaining the vector $\bar{\mathbf{u}}$ is free, since the last row of $\bar{\mathbf{A}}$ is either the last row of $\mathbf{A}$ (which is part of the input) or the zero row. To summarize, we can always find a suitable generator $(\bar{\mathbf{G}}, \bar{\mathbf{H}}, \bar{\mathbf{u}})$ of the augmented matrix $\bar{\mathbf{A}}$ in time $O(\alpha \mathbf{M}(\bar{p}) + \alpha^\omega)$, that is, since $\bar{p} < 2p$ and $\alpha \leq p$,

$$O(\alpha \mathbf{M}(p) + \alpha^{\omega-1}p).$$

Since $\bar{\alpha} \leq \bar{p}$ and $\bar{p}$ is a power of two, Lemma 36 implies that the output $(\ell, \mathbf{Y}, \mathbf{Z}, \mathbf{v})$ of `largest_rec` is obtained in time

$$O(\mathcal{M}_{\text{mat}}''(\bar{p}/\bar{\alpha}, \bar{\alpha}))$$

and satisfies the following: ℓ is the size of the largest leading principal submatrix of $\mathbf{A}$ having generic rank profile; furthermore, denoting this matrix by $\mathbf{A}_\ell$, we have $\mathbf{v}^t$ equal to the first row of $\mathbf{A}_\ell^{-1}$ and $\bar{\mathbf{Y}} = -\mathbf{A}_\ell^{-1}\bar{\mathbf{G}}_\ell$ and $\bar{\mathbf{Z}} = -\mathbf{A}_\ell^{-t}\bar{\mathbf{H}}_\ell$ with $\bar{\mathbf{G}}_\ell$ and $\bar{\mathbf{H}}_\ell$ the $\ell \times \bar{\alpha}$ submatrices consisting of the first ℓ rows of $\bar{\mathbf{G}}$ and $\bar{\mathbf{H}}$. Since $\mathbf{G} \in \mathbb{F}^{m \times \alpha}$ and $\mathbf{H} \in \mathbb{F}^{n \times \alpha}$ are in the upper left corners of $\bar{\mathbf{G}}$ and $\bar{\mathbf{H}}$ and since $\ell \leq \min(m, n)$, we deduce that

$$\bar{\mathbf{G}}_\ell = [\mathbf{G}_\ell \ *] \quad \text{and} \quad \bar{\mathbf{H}}_\ell = [\mathbf{H}_\ell \ *].$$

Consequently, by keeping only the first α columns of each of $\bar{\mathbf{Y}}$ and $\bar{\mathbf{Z}}$, we obtain the matrix pair $(\mathbf{Y}, \mathbf{Z})$ such that

$$\mathbf{Y} = -\mathbf{A}_\ell^{-1}\mathbf{G}_\ell \quad \text{and} \quad \mathbf{Z} = -\mathbf{A}_\ell^{-t}\mathbf{H}_\ell.$$

This concludes the proof of correctness. For the cost, note that the smallest integer power of two greater than or equal to $\bar{p}/\bar{\alpha}$ is less than $4p/\alpha$ and, on the other hand, and that $\bar{\alpha} = O(\alpha)$. Hence $\mathcal{M}_{\text{mat}}''(\bar{p}/\bar{\alpha}, \bar{\alpha}) = O(\mathcal{M}_{\text{mat}}''(p/\alpha, \alpha))$, and the latter expression dominates over the term in $O(\alpha \mathbf{M}(p) + \alpha^{\omega-1}p)$. $\square$

LEMMA 38. *Let $p = \max(m, n)$. Algorithm `lp_inv` works correctly in time*

$$O\left(\mathcal{M}_{\text{mat}}''\left(\frac{p}{\alpha}, \alpha\right)\right).$$

Algorithm $\text{lp_inv}(\mathbf{G}, \mathbf{H}, \mathbf{u})$

Input: $(\mathbf{G}, \mathbf{H}) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha}$ such that $\alpha \leq \min(m, n)$ and $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,0}^t}(\mathbf{A}) = \mathbf{G}\mathbf{H}^t$;
 $\mathbf{u} \in \mathbb{F}^n$ such that $\mathbf{u}^t = \mathbf{e}_{m,m}^t \mathbf{A}$ (the last row of $\mathbf{A}$).

Output: $(r, \mathbf{Y}, \mathbf{Z}, \mathbf{v})$ such that $r = \text{rank}(\mathbf{A})$, $(\mathbf{Y}, \mathbf{Z}) = (-\mathbf{A}_r^{-1} \mathbf{G}_r, \mathbf{A}_r^{-t} \mathbf{H}_r)$, and
 $\mathbf{v} = \mathbf{A}_r^{-t} \mathbf{e}_{r,1}$ (the first row of $\mathbf{A}_r^{-1}$) if $\mathbf{A}$ has generic rank profile,
and $r = \text{"failure"}$ otherwise.

$(\ell, \mathbf{Y}_\ell, \mathbf{Z}_\ell, \mathbf{v}_\ell) := \text{largest}(\mathbf{G}, \mathbf{H}, \mathbf{u})$

$\mathbf{G}_S := \mathbf{G}_2 + \mathbf{A}_{21} \mathbf{Y}_\ell$; $\mathbf{H}_S := \mathbf{H}_2 - \mathbf{A}_{12}^t \mathbf{Z}_\ell$

$\mathbf{u}_S := \mathbf{u}_{22} - \mathbf{A}_{12}^t \mathbf{A}_\ell^{-t} \mathbf{u}_{21}$

$r_0 := \text{the rank of } [\mathbf{G}_S | \mathbf{e}_{m-\ell,1}] [\mathbf{H}_S | \mathbf{u}_S]^t$

if $r_0 = 0$ then

$(r, \mathbf{Y}, \mathbf{Z}, \mathbf{v}) := (\ell, \mathbf{Y}_\ell, \mathbf{Z}_\ell, \mathbf{v}_\ell)$

else

$(r, \mathbf{Y}, \mathbf{Z}, \mathbf{v}) := (\text{"failure"}, *, *, *)$

return $(r, \mathbf{Y}, \mathbf{Z}, \mathbf{v})$.

FIG. 6. Algorithm lp_inv . Here r denotes the (unknown) rank of $\mathbf{A}$, and $\mathbf{A}_r$ is the $r \times r$ matrix consisting of the first r rows and columns of $\mathbf{A}$. Similarly, $\mathbf{G}_r$ and $\mathbf{H}_r$ are the $r \times \alpha$ matrices consisting of the first r rows of $\mathbf{G}$ and $\mathbf{H}$, respectively.

Proof. By Lemma 37, the call to largest yields $(\ell, \mathbf{Y}_\ell, \mathbf{Z}_\ell, \mathbf{v}_\ell)$ such that ℓ is the order of the largest leading principal submatrix of $\mathbf{A}$ having generic rank profile, $\mathbf{Y}_\ell = -\mathbf{A}_\ell^{-1} \mathbf{G}_\ell$, $\mathbf{Z}_\ell = \mathbf{A}_\ell^{-t} \mathbf{H}_\ell$, and $\mathbf{v}_\ell = \mathbf{A}_\ell^{-t} \mathbf{e}_{\ell,1}$. Thanks to the special shape of the matrices $\mathbf{Y}_\ell$ and $\mathbf{Z}_\ell$, we can check as in [16] that the expressions for $\mathbf{G}_S$, $\mathbf{H}_S$, $\mathbf{u}_S$ lead to a $\nabla_{\mathbb{Z}_{m-\ell,1}, \mathbb{Z}_{n-\ell,0}^t}$ -generator of length $\alpha+1$ of the Schur complement $\mathbf{S} = \mathbf{A}_{22} - \mathbf{A}_{21} \mathbf{A}_\ell^{-1} \mathbf{A}_{12}$:

$$\nabla_{\mathbb{Z}_{m-\ell,1}, \mathbb{Z}_{n-\ell,0}^t}(\mathbf{S}) = [\mathbf{G}_S | \mathbf{e}_{m-\ell,1}] [\mathbf{H}_S | \mathbf{u}_S]^t.$$

Now, since both $\nabla_{\mathbb{Z}_{m-\ell,1}, \mathbb{Z}_{n-\ell,0}^t}$ and $\mathbf{A}_\ell$ are invertible,

$$r_0 := \text{rank}(\nabla_{\mathbb{Z}_{m-\ell,1}, \mathbb{Z}_{n-\ell,0}^t}(\mathbf{S})) = 0 \iff \mathbf{S} = 0 \iff \ell = \text{rank}(\mathbf{A}).$$

Correctness then follows, since $\ell = \text{rank}(\mathbf{A})$ if and only if $\mathbf{A}$ has generic rank profile.

To bound the cost, recall first from Lemma 37 that ℓ , $\mathbf{Y}_\ell$, $\mathbf{Z}_\ell$, $\mathbf{v}_\ell$ are computed in time $O(\mathcal{M}_{\text{mat}}''(p/\alpha, \alpha))$. Then we can obtain $\mathbf{G}_S$ in time $O(\mathcal{M}_{\text{mat}}'(p/\alpha, \alpha))$ by evaluating the product $\mathbf{A}_{21} \mathbf{Y}_\ell$ as follows. Defining $\bar{\mathbf{A}}_{21} = \begin{bmatrix} 0 & 0 \\ 0 & \mathbf{A}_{21} \end{bmatrix} \in \mathbb{F}^{m \times n}$ and $\bar{\mathbf{Y}}_\ell = \begin{bmatrix} 0 \\ \mathbf{Y}_\ell \end{bmatrix} \in \mathbb{F}^{n \times \alpha}$, we can deduce from $\mathbf{G}, \mathbf{H}, \mathbf{u}$ a $\nabla_{\mathbb{Z}_{m,1}, \mathbb{Z}_{n,0}^t}$ -generator $(\bar{\mathbf{G}}_{21}, \bar{\mathbf{H}}_{21})$ of length $\alpha+2$ of $\bar{\mathbf{A}}_{21}$ in time $O(\alpha M(p))$. Writing $\bar{\mathbf{A}}_{21} = \bar{\mathbf{A}}_{21}' + \bar{\mathbf{A}}_{21}''$ with $\bar{\mathbf{A}}_{21}'$ of displacement rank $\alpha \leq \min(m, n)$ and $\bar{\mathbf{A}}_{21}''$ of displacement rank 2, we can evaluate $\bar{\mathbf{A}}_{21}' \bar{\mathbf{Y}}_\ell$ in time $\mathcal{C}_{\text{mul}}(\nabla_{\mathbb{Z}_{m,1}, \mathbb{Z}_{n,0}^t}, \alpha, \alpha) + O(\alpha M(p))$, which by Theorem 2 is in $O(\mathcal{M}_{\text{mat}}'(p/\alpha, \alpha) + \alpha M(p))$. It remains to extract $\mathbf{A}_{21} \mathbf{Y}_\ell$ from $\bar{\mathbf{A}}_{21}' \bar{\mathbf{Y}}_\ell = \begin{bmatrix} 0 \\ \mathbf{A}_{21} \mathbf{Y}_\ell \end{bmatrix}$ and to add it to $\mathbf{G}_2$, for an overhead of $O(\alpha p)$. Since $\alpha M(p) = O(\mathcal{M}_{\text{mat}}'(p/\alpha, \alpha))$, we have thus obtained $\mathbf{G}_S$ in time $O(\mathcal{M}_{\text{mat}}'(p/\alpha, \alpha))$. The same cost bound can be derived for the matrix $\mathbf{H}_S$ and, on the other hand, the cost bound $O(\alpha M(p))$ applies to the computation of $\mathbf{u}_S$ and follows from computing a generator of length $O(\alpha)$ of $\mathbf{A}_{12}^t \mathbf{A}_\ell^{-t}$ and then multiplying by and subtracting from a vector. Finally, given $\mathbf{G}_S$, $\mathbf{H}_S$, $\mathbf{u}_S$, the rank r_0 can be deduced in time $O(\alpha^{\omega-1} p)$. The conclusion for the cost then follows from the fact that $\alpha^{\omega-1} p$ and $\mathcal{M}_{\text{mat}}'(p/\alpha, \alpha)$ are both in $O(\mathcal{M}_{\text{mat}}''(p/\alpha, \alpha))$. $\square$

6.3. Generating inverses and solving linear systems. We conclude by showing how to apply algorithm `lp_inv` from the previous section to our initial problems $\text{inv}(\mathcal{L}, \alpha)$ and $\text{solve}(\mathcal{L}, \alpha)$. This corresponds to algorithms `inv` and `solve` given in Figures 7 and 8.

Algorithm `inv`(G, H, S)

Input: $(G, H) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{m \times \alpha}$ such that $\alpha \leq m$ and $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{m,1}^t}(A) = GH^t$,
and a finite subset $S \subset \mathbb{F}$.

Output: if not “failure”, then $(Y, Z) = (-A^{-1}G, A^{-t}H)$ or “A is singular”.

choose the entries of $r_1, r_2 \in \mathbb{F}^{m-1}$ uniformly at random from S

$v_1 := [1 \mid r_1^t]^t$; $v_2 := [1 \mid r_2^t]^t$

$(\tilde{G}, \tilde{H}, \tilde{u}) := \text{precond}(G, H, v_1, v_2)$

$(r, \tilde{Y}, \tilde{Z}, *) := \text{lp_inv}(\tilde{G}, \tilde{H}, \tilde{u})$

if $r \notin \{0, 1, \dots, m\}$ then
return “failure”

else
if $r = m$ then
 $Y := \mathbb{U}(v_2)^t \tilde{Y}^{\mapsto \alpha}$; $Z := \mathbb{U}(v_1)^t \tilde{Z}^{\mapsto \alpha}$
 return (Y, Z)

else
return “A is singular”

FIG. 7. Algorithm `inv`.

THEOREM 39. *Algorithm `inv` works correctly in time $O(\mathcal{M}_{\text{mat}}''(\frac{m}{\alpha}, \alpha))$. It makes $2m-2$ random choices in $\mathbb{F}$ and fails with probability less than $1/2$ if $|S| \geq 2m(m+1)$.*

Proof. By applying Lemma 35 to the case $m = n$, we see that $\tilde{u}^t$ is the last row of $\tilde{A} = \mathbb{U}(v_1)A\mathbb{U}(v_2)^t$ and that $\tilde{G}$ and $\tilde{H}$ are $m \times (\alpha + 4)$ matrices such that

$$\mathbb{Z}_{m,0} \tilde{A} - \tilde{A} \mathbb{Z}_{m,0}^t = \tilde{G} \tilde{H}^t, \quad \tilde{G}^{\mapsto \alpha} = \mathbb{U}(v_1)G, \quad \tilde{H}^{\mapsto \alpha} = \mathbb{U}(v_2)H.$$

If $r \notin \{0, 1, \dots, m\}$ then $r = \text{“failure”}$, and this is what we return. Assume now that $r \in \{0, 1, \dots, m\}$. In this case, preconditioning has ensured that $\tilde{A}$ has generic rank profile. The number r produced by `lp_inv` thus satisfies $r = \text{rank}(\tilde{A}) = \text{rank}(A)$, and A is singular if and only if $r \neq m$. When $r = m$, we have $\tilde{A} = \tilde{A}_r$, so the matrices $\tilde{Y}$ and $\tilde{Z}$ produced by `lp_inv` satisfy $\tilde{Y} = -\tilde{A}^{-1}\tilde{G}$ and $\tilde{Z} = \tilde{A}^{-t}\tilde{H}$. Using the special shape of the first α columns of $\tilde{G}$ and $\tilde{H}$, we conclude that the returned matrices Y and Z are $-A^{-1}G$ and $A^{-t}H$, as wanted. (Note that although it is produced by `lp_inv`, the first row of the inverse of $\tilde{A}_r$ is not needed here and denoted by $*$; we will need it, however, when solving linear systems at the end of this section.)

Applying Lemmas 35 and 38 with $p = m$ shows that the cost of calling `precond` and `lp_inv` is $O(\alpha M(m) + \mathcal{M}_{\text{mat}}''(m/\alpha, \alpha))$; on the other hand, the Toeplitz structure of $\mathbb{U}(v_1)$ and $\mathbb{U}(v_2)$ implies that Y and Z can be deduced in time $O(\alpha M(m))$ from $v_1, v_2, \tilde{Y}$, and $\tilde{Z}$. Writing $\bar{\alpha}$ and $\overline{m/\alpha}$ for the smallest integer powers of two greater than or equal to α and m/α , we have $\mathcal{M}_{\text{mat}}''(m/\alpha, \alpha) \geq \mathcal{M}_{\text{mat}}'(\overline{m/\alpha}, \alpha) \geq \bar{\alpha} \cdot \mathcal{M}_{\text{mat}}(\bar{\alpha} \cdot \overline{m/\alpha}, 1) = \bar{\alpha} M(\bar{\alpha} \cdot \overline{m/\alpha}) \geq \alpha M(m)$, from which the claimed cost bound follows.

Finally, by Lemma 35, the preconditioned matrix $\tilde{\mathbf{A}}$ has generic rank profile with probability $\mathcal{P} \geq 1 - \text{rank}(\mathbf{A}) \cdot (\text{rank}(\mathbf{A}) + 1)/|S|$. Since $\text{rank}(\mathbf{A}) \leq m$, we have $\mathcal{P} \geq 1 - m(m+1)/|S|$, so that $|S| \geq 2m(m+1)$ implies $\mathcal{P} \geq 1/2$. $\square$

For solving $\mathbf{Ax} = \mathbf{b}$, we work on the equivalent (preconditioned) system $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}$ such that $\tilde{\mathbf{A}} = \mathbb{U}(\mathbf{v}_1)\mathbf{A}\mathbb{U}(\mathbf{v}_2)^t$ and $\tilde{\mathbf{b}} = \mathbb{U}(\mathbf{v}_1)\mathbf{b}$, for which any solution $\tilde{\mathbf{x}}$ yields a solution $\mathbf{x} = \mathbb{U}(\mathbf{v}_2)^t\tilde{\mathbf{x}}$. Algorithm solve in Figure 8 uses the following notation: writing as before r for the rank of $\mathbf{A}$, we partition $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{b}}$ into blocks as

$$\tilde{\mathbf{A}} = \begin{bmatrix} \tilde{\mathbf{A}}_r & \tilde{\mathbf{A}}_{12} \\ \tilde{\mathbf{A}}_{21} & \tilde{\mathbf{A}}_{21}\tilde{\mathbf{A}}_r^{-1}\tilde{\mathbf{A}}_{12} \end{bmatrix}, \quad \tilde{\mathbf{b}} = \begin{bmatrix} \tilde{\mathbf{b}}_1 \\ \tilde{\mathbf{b}}_2 \end{bmatrix}$$

with $\mathbf{A}_{12} \in \mathbb{F}^{r \times (n-r)}$, $\tilde{\mathbf{A}}_{21} \in \mathbb{F}^{(m-r) \times r}$, $\tilde{\mathbf{b}}_1 \in \mathbb{F}^r$ and $\tilde{\mathbf{b}}_2 \in \mathbb{F}^{m-r}$.

Algorithm solve($\mathbf{G}, \mathbf{H}, \mathbf{b}, S$)

Input: $(\mathbf{G}, \mathbf{H}) \in \mathbb{F}^{m \times \alpha} \times \mathbb{F}^{n \times \alpha}$ such that $\alpha \leq \min(m, n)$ and $\nabla_{\mathbb{Z}_{m,0}, \mathbb{Z}_{n,1}^t}(\mathbf{A}) = \mathbf{GH}^t$, $\mathbf{b} \in \mathbb{F}^m$, and a finite subset $S \subset \mathbb{F}$.

Output: if not “failure”, then a nontrivial solution $\mathbf{x} \in \mathbb{F}^n$ to $\mathbf{Ax} = \mathbf{b}$ or “no solution exists.”

choose the entries of $\mathbf{r}_1 \in \mathbb{F}^{m-1}$ and $\mathbf{r}_2 \in \mathbb{F}^{n-1}$ uniformly at random from S

$\mathbf{v}_1 := [1 \mid \mathbf{r}_1^t]^t$; $\mathbf{v}_2 := [1 \mid \mathbf{r}_2^t]^t$

$(\tilde{\mathbf{G}}, \tilde{\mathbf{H}}, \tilde{\mathbf{u}}) := \text{precond}(\mathbf{G}, \mathbf{H}, \mathbf{v}_1, \mathbf{v}_2)$

$(r, \tilde{\mathbf{Y}}, \tilde{\mathbf{Z}}, \tilde{\mathbf{v}}) := \text{lp_inv}(\tilde{\mathbf{G}}, \tilde{\mathbf{H}}, \tilde{\mathbf{u}})$

if $r \notin \{0, 1, \dots, m\}$ then
return “failure”

else

$\begin{bmatrix} \tilde{\mathbf{b}}_1 \\ \tilde{\mathbf{b}}_2 \end{bmatrix} := \mathbb{U}(\mathbf{v}_1)\mathbf{b}$

$\tilde{\mathbf{x}}_1 := \tilde{\mathbf{A}}_r^{-1}\tilde{\mathbf{b}}_1$

if $\tilde{\mathbf{A}}_{21}\tilde{\mathbf{x}}_1 = \tilde{\mathbf{b}}_2$ then

$\tilde{\mathbf{x}} := \begin{bmatrix} \tilde{\mathbf{x}}_1 + \tilde{\mathbf{A}}_r^{-1}\tilde{\mathbf{A}}_{12}\mathbf{e}_{n-r,1} \\ -\mathbf{e}_{n-r,1} \end{bmatrix}$

$\mathbf{x} := \mathbb{U}(\mathbf{v}_2)^t\tilde{\mathbf{x}}$

return $\mathbf{x}$

else

return “no solution exists”

FIG. 8. Algorithm solve.

THEOREM 40. *Algorithm solve works correctly in time $O(\mathcal{M}_{\text{mat}}''(\frac{p}{\alpha}, \alpha))$ with $p = \max(m, n)$. It makes $m + n - 2$ random choices in $\mathbb{F}$ and fails with probability less than $1/2$ if $|S| \geq 2q(q+1)$ with $q = \min(m, n)$.*

Proof. Recalling that $\mathbf{Ax} = \mathbf{b}$ is equivalent to $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}$ with $\tilde{\mathbf{A}} = \mathbb{U}(\mathbf{v}_1)\mathbf{A}\mathbb{U}(\mathbf{v}_2)^t$, $\tilde{\mathbf{x}} = \mathbb{U}(\mathbf{v}_2)^{-t}\mathbf{x}$ and $\tilde{\mathbf{b}} = \mathbb{U}(\mathbf{v}_1)\mathbf{b}$, we see that the algorithm begins by generating the matrix $\tilde{\mathbf{A}}$. If $r \in \{0, 1, \dots, m\}$, then $\tilde{\mathbf{A}}$ has generic rank profile and thus $r = \text{rank}(\tilde{\mathbf{A}}) = \text{rank}(\mathbf{A})$. This implies $\tilde{\mathbf{A}}_{22} - \tilde{\mathbf{A}}_{21}\tilde{\mathbf{A}}_r^{-1}\tilde{\mathbf{A}}_{12} = 0$ and, partitioning $\tilde{\mathbf{b}}$ conformally with $\tilde{\mathbf{A}}$,

we deduce that the linear system $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}$ is equivalent to

$$\begin{bmatrix} \tilde{\mathbf{A}}_r & \tilde{\mathbf{A}}_{12} \\ 0 & 0 \end{bmatrix} \tilde{\mathbf{x}} = \begin{bmatrix} \mathbb{I}_r & 0 \\ -\tilde{\mathbf{A}}_{21}\tilde{\mathbf{A}}_r^{-1} & \mathbb{I}_{m-r} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{b}}_1 \\ \tilde{\mathbf{b}}_2 \end{bmatrix}.$$

If $\tilde{\mathbf{b}}_2$ is such that $-\tilde{\mathbf{A}}_{21}\tilde{\mathbf{x}}_1 + \tilde{\mathbf{b}}_2 \neq 0$ for $\tilde{\mathbf{x}}_1 = \tilde{\mathbf{A}}_r^{-1}\tilde{\mathbf{b}}_1$, then no solution exists. Else, it is easily checked that any vector of the form

$$\tilde{\mathbf{x}}(\mathbf{v}) = \begin{bmatrix} \tilde{\mathbf{x}}_1 + \tilde{\mathbf{A}}_r^{-1}\tilde{\mathbf{A}}_{12}\mathbf{v} \\ -\mathbf{v} \end{bmatrix} \quad \text{with } \mathbf{v} \in \mathbb{F}^{n-r}$$

is a solution to $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}$; furthermore, taking $\mathbf{v} = \mathbf{e}_{n-r,1} \neq 0$ when $n-r > 0$ ensures that this solution is nontrivial. Pre-multiplying this solution by $\mathbb{U}(\mathbf{v}_2)^t$ then yields a nontrivial solution to the original system, so correctness follows.

By Lemmas 35 and 38, calling `precond` and `lp_inv` uses $O(\alpha M(p) + \mathcal{M}_{\text{mat}}''(p/\alpha, \alpha))$ operations in $\mathbb{F}$. Then, since $\mathbb{U}(\mathbf{v}_1)$ and $\mathbb{U}(\mathbf{v}_2)$ are Toeplitz matrices, we can deduce $\tilde{\mathbf{b}} = \mathbb{U}(\mathbf{v}_1)\mathbf{b}$ and $\mathbf{x} = \mathbb{U}(\mathbf{v}_2)^t\tilde{\mathbf{x}}$ from $\mathbf{v}_1, \mathbf{v}_2, \mathbf{b}, \tilde{\mathbf{x}}$ in time $O(M(p))$. Finally, given the generators $(\tilde{\mathbf{G}}, \tilde{\mathbf{H}}, \tilde{\mathbf{v}})$ and $(\tilde{\mathbf{Y}}, \tilde{\mathbf{Z}}, \tilde{\mathbf{v}})$ of $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{A}}^{-1}$, we can generate the submatrices $\tilde{\mathbf{A}}_{12}$ and $\tilde{\mathbf{A}}_{21}$ and perform the matrix-vector products $\tilde{\mathbf{A}}_r^{-1}\tilde{\mathbf{b}}_1$, $\tilde{\mathbf{A}}_{21}\tilde{\mathbf{x}}_1$, and $\tilde{\mathbf{A}}_r^{-1}\tilde{\mathbf{A}}_{12}\mathbf{e}_{n-r,1}$ in time $O(\alpha M(p))$. Recalling that $\alpha M(p) \leq \mathcal{M}_{\text{mat}}''(p/\alpha, \alpha)$, we conclude that the total cost is thus in $O(\mathcal{M}_{\text{mat}}''(p/\alpha, \alpha))$.

The probability analysis is the same as for inversion. $\square$

Remark. If one wants to sample uniformly the solution manifold of $\mathbf{A}\mathbf{x} = \mathbf{b}$, then it suffices to replace the unit vector $\mathbf{e}_{n-r,1}$ in algorithm `solve` by a vector $\mathbf{r}_3 \in \mathbb{F}^{n-r}$ whose entries are selected uniformly at random from the subset S . This way of constructing $\mathbf{x}$ corresponds to the technique introduced by Kaltofen and Saunders in [21, Theorem 4], but requires only $n-r$ additional random choices instead of n .

REFERENCES

- [1] A. V. AHO, K. STEIGLITZ, AND J. D. ULLMAN, *Evaluating polynomials at fixed sets of points*, SIAM J. Comput., 4 (1975), pp. 533–539.
- [2] D. BINI AND V. Y. PAN, *Polynomial and Matrix Computations, volume 1: Fundamental Algorithms*, Birkhäuser, 1994.
- [3] R. R. BITMEAD AND B. D. O. ANDERSON, *Asymptotically fast solution of Toeplitz and related systems of linear equations*, Linear Algebra Appl., 34 (1980), pp. 103–116.
- [4] A. BOSTAN, C.-P. JEANNEROD, AND É. SCHOST, *Solving structured linear systems with large displacement rank*, Theoret. Comput. Sci., 407 (2008), pp. 155–181.
- [5] A. BOSTAN, G. LECERF, B. SALVY, É. SCHOST, AND B. WIEBELT, *Complexity issues in bivariate polynomial factorization*, Proceedings of ISSAC’04, ACM, 2004, pp. 42–49.
- [6] A. BOSTAN, G. LECERF, AND É. SCHOST, *Tellegen’s principle into practice*, Proceedings of ISSAC’03, ACM, 2003, pp. 37–44.
- [7] A. BOSTAN AND É. SCHOST, *Polynomial evaluation and interpolation on special sets of points*, J. Complexity, 21 (2005), pp. 420–446.
- [8] P. BÜRGISSER, M. CLAUSEN, AND A. SHOKROLLAHI, *Algebraic Complexity Theory*, Springer, 1997.
- [9] D. G. CANTOR AND E. KALTOFEN, *On fast multiplication of polynomials over arbitrary algebras*, Acta Inform., 28 (1991), pp. 693–701.
- [10] M. F. I. CHOWDHURY, C.-P. JEANNEROD, V. NEIGER, É. SCHOST, AND G. VILLARD, *Faster algorithms for multivariate interpolation with multiplicities and simultaneous polynomial approximations*, IEEE Trans. Inform. Theory, 61 (2015), pp. 2370–2387.
- [11] B. FRIEDLANDER, M. MORF, T. KAILATH, AND L. LJUNG, *New inversion formulas for matrices classified in terms of their distance from Toeplitz matrices*, Linear Algebra Appl., 27 (1979), pp. 31–60.

- [12] J. VON ZUR GATHEN AND J. GERHARD, *Modern computer algebra*, Cambridge University Press, third ed., 2013.
- [13] I. GOHBERG AND V. OLSHEVSKY, *Complexity of multiplication with vectors for structured matrices*, Linear Algebra Appl., 202 (1994), pp. 163–192.
- [14] I. GOHBERG AND V. OLSHEVSKY, *Fast algorithms with preprocessing for matrix-vector multiplication problems*, J. Complexity, 10 (1994), pp. 411–427.
- [15] G. HANROT, M. QUERCIA, AND P. ZIMMERMANN, *The middle product algorithm. I*, Appl. Algebra Engrg. Comm. Comput., 14 (2004), pp. 415–438.
- [16] C.-P. JEANNEROD AND C. MOUILLERON, *Computing specified generators of structured matrix inverses*, Proceedings of ISSAC’10, ACM, 2010, pp. 281–288.
- [17] T. KAILATH, S. Y. KUNG, AND M. MORF, *Displacement ranks of a matrix*, Bull. Amer. Math. Soc. (N.S.), 1 (1979), pp. 769–773.
- [18] T. KAILATH, S. Y. KUNG, AND M. MORF, *Displacement ranks of matrices and linear equations*, J. Math. Anal. Appl., 68 (1979), pp. 395–407.
- [19] E. KALTOFEN, *Asymptotically fast solution of Toeplitz-like singular linear systems*, Proceedings of ISSAC’94, ACM, 1994, pp. 297–304.
- [20] E. KALTOFEN, *Analysis of Coppersmith’s block Wiedemann algorithm for the parallel solution of sparse linear systems*, Math. Comp., 64 (1995), pp. 777–806.
- [21] E. KALTOFEN AND D. SAUNDERS, *On Wiedemann’s method of solving sparse linear systems*, in AAEECC-9, vol. 539 of Lecture Notes in Computer Science, Springer, 1991, pp. 29–38.
- [22] P. LANCASTER AND M. TISMENETSKY, *The Theory of Matrices, Second Edition*, Computer Science and Applied Mathematics, Academic Press, 1985.
- [23] F. LE GALL, *Powers of tensors and fast matrix multiplication*, in Proceedings of ISSAC’14, ACM, 2014, pp. 296–303.
- [24] M. MORF, *Fast Algorithms for Multivariable Systems*, PhD thesis, Dept. of Electrical Engineering, Stanford University, Stanford, 1974.
- [25] M. MORF, *Doubling algorithms for Toeplitz and related equations*, in Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, 1980, pp. 954–959.
- [26] V. OLSHEVSKY AND V. PAN, *A unified superfast algorithm for boundary rational tangential interpolation problems and for inversion and factorization of dense structured matrices*, in Proc. 39th IEEE FOCS, 1998, pp. 192–201.
- [27] V. OLSHEVSKY AND A. SHOKROLLAHI, *A unified superfast algorithm for confluent tangential interpolation problem and for structured matrices*, in Advanced Signal Processing Algorithms, Architectures, and Implementations, ASPAAPIX, SPIE, 1999, pp. 312–323.
- [28] V. OLSHEVSKY AND A. SHOKROLLAHI, *Matrix-vector product for confluent Cauchy-like matrices with application to confluent rational interpolation*, in STOC’00, ACM, 2000, pp. 573–581.
- [29] V. Y. PAN, *Trilinear aggregating with implicit canceling for a new acceleration of matrix multiplication*, Comp. & Maths. with Appls., 8 (1982), pp. 23–34.
- [30] V. Y. PAN, *On computations with dense structured matrices*, Math. Comp., 55 (1990), pp. 179–190.
- [31] V. Y. PAN, *A unified superfast divide-and-conquer algorithm for structured matrices*. MSRI Preprint 1999-033, Mathematical Sciences Research Institute, Berkeley, CA, April 1999.
- [32] V. Y. PAN, *Nearly optimal computations with structured matrices*, in SODA’00, ACM, 2000, pp. 953–962.
- [33] V. Y. PAN, *Structured Matrices and Polynomials*, Birkhäuser Boston Inc., 2001.
- [34] V. Y. PAN, *Transformations of matrix structures work again*, Linear Algebra Appl., 465 (2015), pp. 107–138.
- [35] V. Y. PAN AND X. WANG, *Inversion of displacement operators*, SIAM J. Matrix Anal. Appl., 24 (2003), pp. 660–677.
- [36] V. Y. PAN AND A. ZHENG, *Superfast algorithms for Cauchy-like matrix computations and extensions*, Linear Algebra Appl., 310 (2000), pp. 83–108.
- [37] A. SCHÖNHAGE, *Schnelle Multiplikation von Polynomen über Körpern der Charakteristik 2*, Acta Inform., 7 (1977), pp. 395–398.
- [38] A. SCHÖNHAGE AND V. STRASSEN, *Schnelle Multiplikation großer Zahlen*, Computing, 7 (1971), pp. 281–292.
- [39] I. S. SERGEEV, *Fast algorithms for elementary operations with complex power series*, Diskret. Mat., 22 (2010), pp. 17–49.
- [40] V. STRASSEN, *Gaussian elimination is not optimal*, Numer. Math., 13 (1969), pp. 354–356.

Appendix A. Proof that $\mathcal{L}(A)$ and $\mathcal{L}'(A^{-1})$ have the same rank for $\mathcal{L}$ and $\mathcal{L}'$ as in (1). Consider first the Sylvester case, where $\mathcal{L}(A) = MA - AN$ and $\mathcal{L}'(A^{-1}) = NA^{-1} - A^{-1}M$. If $\mathcal{L}(A) = GH^t$, then pre- and post-multiplying both sides of this equality by A^{-1} gives $A^{-1}M - NA^{-1} = A^{-1}GH^tA^{-1}$, so that $\mathcal{L}'(A^{-1})$ equals $-A^{-1}G(A^{-t}H)^t$ and thus has the same rank as $\mathcal{L}(A)$.

Consider now the Stein case, where $\mathcal{L}(A) = A - MAN$ and $\mathcal{L}'(A^{-1}) = A^{-1} - NA^{-1}M$. Defining the matrix

$$B = \begin{bmatrix} A & M \\ N & A^{-1} \end{bmatrix},$$

we have

$$\begin{bmatrix} \mathbb{I} & -MA \\ & \mathbb{I} \end{bmatrix} B \begin{bmatrix} \mathbb{I} & \\ -AN & \mathbb{I} \end{bmatrix} = \text{diag}(A - MAN, A^{-1})$$

and

$$\begin{bmatrix} \mathbb{I} & \\ -NA^{-1} & \mathbb{I} \end{bmatrix} B \begin{bmatrix} \mathbb{I} & -A^{-1}M \\ & \mathbb{I} \end{bmatrix} = \text{diag}(A, A^{-1} - NA^{-1}M).$$

Since the four matrices applied to B are invertible, we deduce that

$$\text{rank}(B) = \text{rank}(A - MAN) + \text{rank}(A^{-1}) = \text{rank}(A) + \text{rank}(A^{-1} - NA^{-1}M).$$

Using $\text{rank}(A) = \text{rank}(A^{-1})$, we see again that $\mathcal{L}(A)$ and $\mathcal{L}'(A^{-1})$ have the same rank.

Appendix B. Proof of Lemma 18. Fix i and j . For all $\ell \geq 1$, [35, Theorem 4.7] gives

$$A'_{i,j} - M_{P_i}^\ell A'_{i,j} (M_{Q_j}^t)^\ell = \sum_{k \leq \alpha} \mathbb{K}(M_{P_i}, g_{i,k}, \ell) \mathbb{K}(M_{Q_j}, h_{j,k}, \ell)^t$$

with $\mathbb{K}(M_{P_i}, g_{i,k}, \ell)$ and $\mathbb{K}(M_{Q_j}, h_{j,k}, \ell)$ as in the proof of Lemma 15. Writing $Q_j = q_{j,0} + \dots + q_{j,n_j} x^{n_j}$, and multiplying the former equality on the left by $q_{j,\ell} M_{P_i}^{n_j - \ell}$, we get, for $1 \leq \ell \leq n_j$,

$$\begin{aligned} q_{j,\ell} M_{P_i}^{n_j - \ell} A'_{i,j} - q_{j,\ell} M_{P_i}^{n_j} A'_{i,j} (M_{Q_j}^t)^\ell \\ = \sum_{k \leq \alpha} M_{P_i}^{n_j - \ell} \mathbb{K}(M_{P_i}, g_{i,k}, \ell) q_{j,\ell} \mathbb{K}(M_{Q_j}, h_{j,k}, \ell)^t \\ = \sum_{k \leq \alpha} \mathbb{K}(M_{P_i}, g_{i,k}, n_j) \mathbb{J}_{n_j} q_{j,\ell} \mathbb{J}_{\ell, n_j} \mathbb{K}(M_{Q_j}, h_{j,k}, n_j)^t, \end{aligned}$$

since the last ℓ columns of $\mathbb{K}(M_{P_i}, g_{i,k}, n_j)$ are precisely $M_{P_i}^{n_j - \ell} \mathbb{K}(M_{P_i}, g_{i,k}, \ell)$; note that the equality also holds for $\ell = 0$. Summing over all $\ell = 0, \dots, n_j$, and using the fact that $Q_j(M_{Q_j}^t) = 0$, we deduce

$$\widetilde{Q}_j(M_{P_i}) A'_{i,j} = \sum_{k \leq \alpha} \mathbb{K}(M_{P_i}, g_{i,k}, n_j) \mathbb{J}_{n_j} \mathbb{Y}_{Q_j} \mathbb{K}(M_{Q_j}, h_{j,k}, n_j)^t.$$

The rest of the proof now follows exactly that of Lemma 15.