Algorithms for structured linear systems solving and their implementation

Seung Gyu Hyun, Romain Lebreton, Éric Schost

► To cite this version:

Seung Gyu Hyun, Romain Lebreton, Éric Schost. Algorithms for structured linear systems solving and their implementation. ISSAC 2017 - 42nd International Symposium on Symbolic and Algebraic Computation, Jul 2017, Kaiserslautern, Germany. pp.205-212, 10.1145/3087604.3087659 . lirmm-01484831

HAL Id: lirmm-01484831

<https://hal-lirmm.ccsd.cnrs.fr/lirmm-01484831>

Submitted on 7 Mar 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.

Algorithms for structured linear systems solving and their implementation

Seung Gyu Hyun
Cheriton School of Computer Science
University of Waterloo
sghyun@edu.uwaterloo.ca

Romain Lebreton
LIRMM
Université de Montpellier
romain.lebreton@lirmm.fr

Éric Schost
Cheriton School of Computer Science
University of Waterloo
eschost@uwaterloo.ca

ABSTRACT

There exists a vast literature dedicated to algorithms for structured matrices, but relatively few descriptions of actual implementations and their practical performance. In this paper, we consider the problem of solving Cauchy-like systems, and its application to mosaic Toeplitz systems, in two contexts: first in the unit cost model (which is a good model for computations over finite fields), then over $\mathbb{Q}$. We introduce new variants of previous algorithms and describe an implementation of these techniques and its practical behavior. We pay a special attention to particular cases such as the computation of algebraic approximants.

1 INTRODUCTION

In this paper, we discuss linear algebra algorithms inspired by the following kind of question: given polynomials such as

$$\begin{cases} t_0 &= 8x^4 - 8x^2 + 1 \\ t_1 &= 16x^5 - 20x^3 + 5x \\ t_2 &= 32x^6 - 48x^4 + 18x^2 - 1, \end{cases}$$

find that the relation $t_0 - 2xt_1 + t_2 = 0$ holds (this is the recurrence between Chebyshev polynomials of the first kind).

More precisely, given input polynomials $(t_0, \dots, t_{s-1})$ over a field $\mathbb{K}$, together with integers $(n_0, \dots, n_{s-1})$ and σ , the *Hermite-Padé* approximation problem asks to compute polynomials $(p_0, \dots, p_{s-1})$, not all zero, such that $\deg(p_i) < n_i$ holds for all i , and such that we have $p_0 t_0 + \dots + p_{s-1} t_{s-1} = O(x^\sigma)$. There exist numerous applications to this type of question, very important particular cases being *algebraic approximants* (with $t_i = f^i$, for some given f) or *differential approximants* (with $t_i = d^i f / dx^i$, for some given f); see for instance [6, Chapitre 7].

Expressed in the canonical monomial bases, the matrix of a Hermite-Padé problem has size $\sigma \times (n_0 + \dots + n_{s-1})$, and consists of s lower triangular Toeplitz blocks. More generally, our goal in this paper is to compute efficiently elements in the kernel of *mosaic Toeplitz* matrices [23], that is, $m \times n$ matrices $T = (T_{i,j})_{1 \leq i \leq p, 1 \leq j \leq q}$ with a $p \times q$ block structure, each of the pq blocks $T_{i,j}$ being Toeplitz.

An $m \times n$ Toeplitz matrix $T = (t_{i-j})_{1 \leq i \leq m, 1 \leq j \leq n}$ can be succinctly represented by the polynomial $P_T = t_{-n+1} + t_{-n+2}x + \dots + t_{m-1}x^{m+n-2}$; multiplication of T by a vector $b = [b_0 \dots b_{n-1}]^t$ amounts to computing $P_T P_b \bmod x^{m+n-1}$ and keeping the coefficients of degrees $n-1, \dots, m+n-2$. More generally, a mosaic Toeplitz $T = (T_{i,j})_{1 \leq i \leq p, 1 \leq j \leq q}$ can be described by a sequence of pq polynomials $\mathcal{P} = (P_{T_{i,j}})_{1 \leq i \leq p, 1 \leq j \leq q}$, together with the sequences

$I = (m_1, \dots, m_p)$ and $J = (n_1, \dots, n_q)$ giving the row-sizes and column-sizes of the blocks. Then, our main problem is as follows.

PROBLEM A. *Given $\mathcal{P}$ and integers I, J as above, defining a mosaic Toeplitz matrix T , find a non-zero vector in the kernel of T .*

We consider two situations, first over an arbitrary field $\mathbb{K}$, counting all operations in $\mathbb{K}$ at unit cost, then over $\mathbb{Q}$, taking bit-size into account. In all this paper, we closely follow the existing formalism of *structured matrix computations* developed in previous work by Morf [34], Bitmead-Anderson [3], Kailath and co-authors [28, 29], Pan [37, 38], Kaltofen [30], Cardinal [14], etc.

We complement the existing literature as follows. First, we define a class of Cauchy-like matrices (see definitions below) for which the matrix-vector product is faster by a constant factor than previous designs. Next, we show how Pan's technique of "multiplicative transformation" of operators results in matrices that have generic rank profile (with high probability), so that further regularization is not needed in general. We then describe an improved iterative algorithm (of quadratic complexity with respect to the matrix size), that makes use of fast matrix multiplication; finally, for matrices defined over $\mathbb{Q}$, we introduce a divide-and-conquer algorithm as an alternative to Newton iteration, and we show how it can be improved in the case of algebraic approximation.

Another contribution of this paper is a discussion of the design and practical performance of a C++ implementation of these algorithms. To our knowledge, only a few papers address these methods from the practical viewpoint. An early reference is [42], which concludes that the divide-and-conquer (MBA) algorithm for solving Toeplitz matrices in quasi-linear time would require matrices of size 10^6 to break even with quadratic-time algorithms; a more recent article [24] estimates the crossover point to be around 8000.

Our experiments first consider computations over finite fields. We assess the practical impact of fast matrix multiplication for structured matrix algorithms (such as in the new algorithm mentioned above, or those in [8, 9]), and we estimate for what matrix size quasi-linear algorithms become effective (with much lower crossover points than above). In the context of computations over $\mathbb{Q}$, we show that our divide-and-conquer algorithm outperforms Newton iteration consistently, and we demonstrate significant speed-ups for the case of algebraic approximants.

We start in Section 2 with a review of basic results on structured matrices, including a discussion of transformation of operators and regularization. Section 3 describes algorithms applicable in a unit cost model, with in particular a new algorithm that uses fast matrix multiplication, and a discussion of the practical performance of these algorithms. Finally, Section 4 presents lifting algorithms for structured matrices, and the corresponding implementation.

2 BASIC RESULTS

This section reviews background material on structured matrices; for a more comprehensive treatment, we refer the reader to [39].

Overview. Developed in [28], the *displacement operator* approach associates to a matrix A its displacement $\nabla(A)$, that is, the image of A under a *displacement operator* ∇ . Then, we say that A is structured with respect to ∇ if $\nabla(A)$ has a small rank compared to its size; the rank of $\nabla(A)$ is called the *displacement rank* of A with respect to ∇ . A prominent example is the family of so-called *Toeplitz-like* matrices, which are structured for the Toeplitz displacement operator

$$\phi : A \mapsto (Z_m A - A Z_n) = (A \downarrow) - (A \leftarrow)$$

where the $n \times n$ lower shift matrix Z_n is the matrix with ones below the diagonal. The displacement rank of a Toeplitz matrix for this operator is at most two; the displacement rank of a mosaic Toeplitz with a $p \times q$ block structure is at most $p + q$.

The key idea of most algorithms for structured matrices is summarized by Pan's motto [39]: compress, operate, decompress. Indeed, for A of size $m \times n$ over a field $\mathbb{K}$, if $\nabla(A)$ has rank α , it can be represented using few elements through ∇ -generators, that is, two matrices (G, H) in $\mathbb{K}^{m \times \alpha} \times \mathbb{K}^{n \times \alpha}$, with $\nabla(A) = GH^t$; α is the *length* of the generators. The main idea behind algorithms for structured matrices is to use generators as a compact data structure, involving $\alpha(m + n)$ field elements instead of mn .

Cauchy-like matrices. Beyond the Toeplitz structure (and the directly related Hankel one), two other important cases are the so-called Vandermonde and Cauchy structures. While the case of Toeplitz-like matrices was the first one to be studied in detail, we will actually focus on Cauchy-like matrices, as we will see that this particular structure is quite convenient to work with.

For a sequence $u = (u_1, \dots, u_m)$ in $\mathbb{K}^m$, let $D_u \in \mathbb{K}^{m \times m}$ be the diagonal matrix with entries $u_1, \dots, u_m$. Then, given u as above and v in $\mathbb{K}^n$, we will consider the operator $\nabla_{u,v} : A \in \mathbb{K}^{m \times n} \mapsto D_u A - A D_v$; *Cauchy-like* matrices (with respect to the choice of u and v) are those matrices A for which $\nabla_{u,v}(A)$ has small rank.

Let u, v be given and suppose that $u_i \neq v_j$ holds for all i, j . Then, the operator $\nabla_{u,v}$ is invertible: given $\nabla_{u,v}$ -generators (G, H) of length α for A , we can reconstruct A as

$$A = \sum_{i=1}^{\alpha} D_{g_i} C_{u,v} D_{h_i}, \quad C_{u,v} = \left[\frac{1}{u_i - v_j} \right]_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}}, \quad (1)$$

where g_i and h_i are the i th columns of respectively G and H , and matrix $C_{u,v}$ is known as a *Cauchy matrix*. Remark that we can equivalently rewrite A as

$$A = (GH^t) \odot C_{u,v}, \quad (2)$$

where $\odot$ denotes the entrywise product.

We will have to handle submatrices of A through their generators. The fact that D_u and D_v are diagonal matrices makes this easy (this is one of the aspects in which the Cauchy structure behaves more simply than the Toeplitz one). Suppose that (G, H) are generators for A , with respect to the operator $\nabla_{u,v}$, and let $u_I = (u_i)_{i \in I}$ and $v_J = (v_j)_{j \in J}$ be subsequences of respectively u and v , corresponding to entries of indices I and J . Let $A_{I,J}$ be the submatrix of A obtained by keeping rows and columns of indices respectively in I and J , and let (G_I, H_J) be the matrices obtained from (G, H) by

respectively keeping rows of G of indices in I , and rows of H of indices in J . Then, (G_I, H_J) is a ∇_{u_I, v_J} -generator for $A_{I,J}$.

Another useful property relates to inverses of Cauchy-like matrices. If a matrix $A \in \mathbb{K}^{n \times n}$ is invertible, and is structured with respect to an operator $\nabla_{u,v}$, its inverse is structured with respect to $\nabla_{v,u}$: if $D_u A - A D_v = GH^t$, one easily deduces that $D_v A^{-1} - A^{-1} D_u = -(A^{-1} G)(A^{-1} H)^t$, where A^{-t} is a shorthand for $(A^{-1})^t$.

Algorithms. In most of the paper, we give costs in an algebraic model, counting base field operations at unit cost (in Section 4, we work over $\mathbb{Q}$ and use a boolean model).

We let $\mathcal{M}$ be such that over any ring, polynomials of degree at most d can be multiplied in $\mathcal{M}(d)$ base ring operations; we also assume that the super-linearity assumptions of [21, Chapter 8] hold. Using the Cantor-Kaltofen algorithm [13], we can take $\mathcal{M}(d) \in O(d \log(d) \log \log(d))$. We let ω be a feasible exponent for linear algebra, in the sense that matrices of size n can be multiplied in $O(n^\omega)$ base ring operations over any ring; the best bound to date is $\omega < 2.38$ [16, 31]. We will have to compute rank and rank profile of dense matrices; the cost reduces to that of matrix multiplication [25]. The notation $O^{\sim}(\cdot)$ indicates that we omit polylogarithmic terms.

Matrix-vector multiplication with $C_{u,v}$ reduces to degree n polynomial interpolation at the points v and evaluation at the points u . Using fast polynomial evaluation and interpolation, this can be done in time $O(\mathcal{M}(v) \log(v))$, with $v = \max(m, n)$; thus, we can multiply matrix A of (1) by a vector in time $O(\alpha \mathcal{M}(v) \log(v)) \subset O^{\sim}(\alpha v)$. In [39, Theorem 4.7.3], Pan shows that if the entries of both u and v are in *geometric progression*, one can reduce the cost of the matrix-vector multiplication by $C_{u,v}$ to $O(\mathcal{M}(v))$, since polynomial evaluation or interpolation at n points in geometric progression can be done in time $O(\mathcal{M}(n))$ [4, 11]; then, multiplication by A as above takes time $O(\alpha \mathcal{M}(v))$.

Remark 2.1. We propose here a refinement of this idea, that allows us to save a constant factor in runtime: we require that u and v be geometric progressions with the *same ratio* τ . Then, the Cauchy matrix $C_{u,v}$ has entries $1/(u_i - v_j) = 1/(u_1 \tau^{i-1} - v_1 \tau^{j-1})$, so it can be factored as

$$C_{u,v} = D_{\tau^{-1}} \left[\frac{1}{u_i - v_j \tau^{j-i}} \right]_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}},$$

where $D_{\tau^{-1}}$ is diagonal with entries $(1, \tau^{-1}, \tau^{-2}, \dots, \tau^{1-m})$, and where the right-hand matrix is Toeplitz. In the reconstruction formula (1), the diagonal matrix $D_{\tau^{-1}}$ commutes with all matrices D_{g_i} , so we can take it out of the sum. Hence, we replaced α evaluations / interpolations at geometric progressions by α product by Toeplitz matrices, each of which can be done in a single polynomial multiplication. The cost for a matrix-vector product by A remains $O(\alpha \mathcal{M}(v))$, but the constant in the big- O is lower: for $m = n$, using middle product techniques [10, 22], the cost goes down from $3\alpha \mathcal{M}(v) + O(\alpha v)$ to $\alpha \mathcal{M}(v) + O(\alpha v)$.

If one needs to multiply A by *several* vectors, further improvements are possible: we mention without giving details an algorithm from [8], that itself follows [9], and which makes it possible to multiply A by α vectors in time $O(\alpha^{\omega-1} \mathcal{M}(v))$ instead of $O(\alpha^2 \mathcal{M}(v))$, by reduction to a sequence of polynomial matrix multiplications.

Reduction from mosaic Toeplitz to Cauchy structure. Our primary interest lies in mosaic Toeplitz matrices. An important insight of Pan [37] shows that one can reduce questions about Toeplitz-like

matrices to ones about Cauchy-like matrices (and conversely, if one wishes to), for a moderate cost overhead. In this paragraph, we give details on this transformation, following [39, Chapter 4.8].

To a vector u in $\mathbb{K}^m$, let us associate the corresponding Vandermonde matrices $V_u = [u_i^{j-1}]_{1 \leq i, j \leq m}$, $W_u = [u_j^{m-i}]_{1 \leq i, j \leq m}$. For u as above, v in $\mathbb{K}^n$ and T in $\mathbb{K}^{m \times n}$, we let $A = V_u T W_v$; we now prove that if T is Toeplitz-like, A is Cauchy-like.

The $\nabla_{u,v}$ -generators of A depend on the ϕ -generators of T , so let us start with those. If $T = (T_{i,j})_{1 \leq i \leq p, 1 \leq j \leq q}$ is mosaic Toeplitz of size $m \times n$, with $T_{i,j}$ of size $m_i \times n_j$, we define the following.

For $1 \leq i \leq p$, let $T_{i,*}$ be the block matrix $[T_{i,1} \cdots T_{i,q}]$. Let η_i be the first row of $T_{i,*}$, shifted left once, let η'_i be the last row of $T_{i,*}$, and finally let $h_i = (\eta'_{i-1} - \eta_i)$ (for $i = 1$, η'_0 is the zero vector); this is a row-vector of size n . If $e_{i,n}$ denotes the row-vector of dimension n with only a one at index i , we also let g_i be the transpose of $e_{m_1+\dots+m_{i-1}+1,m}$. For $1 \leq j \leq q$, define $T_{*,j}$ similarly to $T_{i,*}$; λ_j and λ'_j are its first and last columns, with now a downward shift for λ'_j ; in addition, their entry of index $m_1+\dots+m_j$ are set to zero. We define $g_{p+j} = (\lambda'_j - \lambda_{j+1})$ (λ_{q+1} is set to zero) and let $h_{p+j} = e_{n_1+\dots+n_j,n}$. With these definitions, the matrices $G = [g_1 \cdots g_{p+q}]$ and $H = [h_1^t \cdots h_{p+q}^t]$ are ϕ -generators of T .

To obtain A , we conjugate T by Vandermonde matrices, since $D_u V_u - V_u Z_m$ and $W_v D_v - Z_n W_v$ have small rank; explicitly, they are equal to respectively $(u^m)^t e_{m,m}$ and $(e_{1,n})^t v^n$, where $u^m = [u_1^m \cdots u_n^m]$, and similarly for v^n . As a consequence, since $\phi(T) = GH^t$, we have

$$\nabla_{u,v}(A) = V_u \phi(T) W_v + (u^m)^t e_{m,m} T W_v - V_u T (e_{1,n})^t v^n. \quad (3)$$

This motivates the following definitions. For $1 \leq i \leq p$, let $h'_i = h_i W_v$ and also let g'_i be the column of V_u of index $m_1 + \dots + m_{i-1} + 1$. For $1 \leq j \leq q$, define $g'_{p+j} = V_u g_{p+j}$ and also let h'_{p+j} be the row of index $n_1 + \dots + n_j$ in W_v . We define $g_{p+q+1} = (u^m)^t$, and h_{p+q+1} as the last row of $T W_v$; we also define g_{p+q+2} as the first column of $-V_u T$, and $h_{p+q+2} = v^n$.

By Eq. (3), $G' = [g'_1 \cdots g'_{p+q+2}]$ and $H' = [h'_1 \cdots h'_{p+q+2}]$ are $\nabla_{u,v}$ -generators of A of length $p+q+2$ (whereas T has ϕ -generators of length $p+q$). The bottleneck in the computation of G' and H' are p left-products by W_v and q products by V_u . If all entries of u and v are in geometric progression, it takes time $O(p \mathcal{M}(n) + q \mathcal{M}(m))$.

Regularization. In our algorithms for Cauchy-like matrices, we will assume that the input matrix has *generic rank profile*, that is, that its leading principal minors of size up to its rank are invertible. Regularization for structured matrices was introduced for this purpose by Kaltofen [30], for the Toeplitz structure. In our Cauchy context, one could apply to A as defined above the regularization procedure from [39, Section 5.6], which consists in replacing A by $A' = D_x C_{a,u} A C_{v,b} D_y$, for some new vectors $x, a \in \mathbb{K}^m$ and $y, b \in \mathbb{K}^n$. Theorem 5.6.2 in [39] shows that if a, u consist of $2m$ distinct scalars, and b, v consist of $2n$ distinct scalars, then there exists a non-zero polynomial Δ in the entries of x and y , of degree at most $\mu = \min(m, n)$ in each block of variables, such that the non-vanishing of Δ implies that A' has generic rank profile (that theorem is stated for square matrices, but the result holds in the rectangular case as well). The downside of this construction is that it requires to compute a pair of generators of length $p+q+4$ for A' , involving the multiplication of G' and H' by $C_{a,u}$ and $C_{v,b}$.

We now prove a useful property, involving only $A = V_u T W_v$ as above. For the rest of this paragraph, we assume that the entries of u and v are indeterminates over $\mathbb{K}$, and we show that A has generic rank profile; this will imply the same property for a generic choice of u, v with entries in $\mathbb{K}$. This construction is clearly favorable over the one above, since it involves no extra computation on A .

Following the proof of [39, Theorem 5.6.2], we can use the Cauchy-Binet formula to express the minors of A in terms of those of V_u, T and W_v . Let $i \leq \text{rank}(T)$ and let $I = \{1, \dots, i\}$. The determinant δ_i of the i th leading principal minor of A is the sum over all $J = \{j_1, \dots, j_i\} \subset \{1, \dots, m\}$ and $K = \{k_1, \dots, k_i\} \subset \{1, \dots, n\}$ of $\alpha_{I,J} \beta_{J,K} \gamma_{K,I}$, where $\alpha_{I,J}$ is the determinant of $(V_u)_{I,J}$, $\beta_{J,K}$ is the determinant of $T_{J,K}$, and $\gamma_{K,I}$ is the determinant of $(W_v)_{K,I}$.

Let $<$ be the monomial ordering on the variables $u_1, \dots, u_i, v_1, \dots, v_i$ that first sorts monomials using the lexicographic order on $u_1, \dots, u_i$ with $u_1 < \dots < u_i$, and then breaks the ties using the lexicographic order on $v_1, \dots, v_i$ with $v_1 > \dots > v_i$. If $J = \{j_1, \dots, j_i\}$ with $j_1 < \dots < j_i$ then the leading monomial (denoted $\text{lm}(\cdot)$) of $\alpha_{I,J}$ is $u^J := u_1^{j_1-1} \cdots u_i^{j_i-1}$. Similarly if $K = \{k_1, \dots, k_i\}$ with $k_1 < \dots < k_i$, then $v^K := v_1^{n-k_1} \cdots v_i^{n-k_i} = \text{lm}(\gamma_{K,I})$. Since T is of rank greater or equal to i , at least one of its i th minor $\beta_{J,K}$ does not vanish. Let $J_{\max}, K_{\max}$ be the pair of subsets that maximizes $u^J v^{K_{\max}}$ among those for which $\beta_{J,K} \neq 0$. Then we must have $\text{lm}(\det(A_{I,I})) = u^{J_{\max}} v^{K_{\max}}$, which shows that $\det(A_{I,I})$ is non-zero. The partial degree of $\det(A_{I,I})$ in any variable is at most $\max(n, m)$.

3 OVER AN ABSTRACT FIELD

In this section, we work over a field $\mathbb{K}$, and we explain how to solve Problem A by using Pan's reduction to the Cauchy-like Problem B below. Even though our goal is to solve a linear system, the algorithms do slightly more: they compute the inverse of a given matrix (or of a maximal minor thereof); this is similar to what happens for dense matrices, where it is not known how to solve linear systems with an exponent better than ω . Then, the main question of this section is the following (the assumptions on the vectors u, v are slightly stronger than the one required for $\nabla_{u,v}$ to be invertible).

PROBLEM B. Consider $u = (u_1, \dots, u_m)$ and $v = (v_1, \dots, v_n)$, such that (u, v) has $m+n$ distinct entries. Given $\nabla_{u,v}$ -generators of length α for A in $\mathbb{K}^{m \times n}$, with $\alpha \leq \min(m, n)$, do the following.

If A does not have generic rank profile, raise an error; else, return $\nabla_{v',u'}$ -generators for the inverse of the leading principal minor A_r of A of order r , with $v' = (v_1, \dots, v_r)$, $u' = (u_1, \dots, u_r)$ and $r = \text{rank}(A)$.

To solve an instance of Problem A, with input matrix T , we apply the transformation of the previous section, to obtain a Cauchy-like matrix $A = V_u T W_v$. We compute generators of A_r^{-1} , where A_r is the maximal leading minor of A , and we return a vector b with $b = W_v \begin{bmatrix} -A_r^{-1} B_c \\ c \end{bmatrix}$, where $A_{r,*} = [A_r \quad B]$ and $c \in \mathbb{K}^{n-r}$ random.

There exist two major classes of algorithms for handling Problem B, iterative ones, of cost $\Theta(mn)$ (for fixed α), and divide-and-conquer algorithms, of quasi-linear cost in $m+n$. We stress that having a fast quadratic-time algorithm is actually crucial in practice: as is the case for the Half-GCD, fast linear algebra algorithms, etc, the divide-and-conquer algorithm will fall back on the iterative one for input sizes under a certain threshold, and the performance of the latter will be an important factor in the overall runtime.

Iterative algorithms that solve a size n Toeplitz system in time $O(n^2)$ have been known for decades [18, 33, 45]; extensions to structured matrices were later given, as for instance in [27]. After a section of preliminary results, we give in Section 3.2 an algorithm inspired by [35, Algorithm 4], for the specific form of our Problem B. In this reference, Moulleron sketches an algorithm that solves Problem B in time $O(\alpha n^2)$, in the case where $m = n$ and A is invertible (but without the rank profile assumption); he credits the origin of this algorithm to Kailath [29, §1.10], who dealt with symmetric matrices. Our algorithm follows the same pattern, but reduces the cost to $O(\alpha^{\omega-2}mn)$.

In Section 3.3, we review divide-and-conquer techniques. Kaltofen [30] gave a divide-and-conquer algorithm that solves the analogue of Problem B for Toeplitz-like matrices, lifting assumptions of (strong) non-singularity needed in the original Morf and Bitmead-Anderson algorithm [3, 34]; a generalization to the Cauchy case is in [14, 40]. A further improvement due to Jeannerod and Moulleron [26], following [14], allows one to bypass costly compression stages that are needed in Kaltofen's algorithm and its extensions, by predicting the shape of the generators we have to compute. For the case of square Cauchy-like matrices of size n , this results in an algorithm of cost $O(\alpha^2 \mathcal{M}(n) \log(n)^2)$, but we will point out that better estimates are available by choosing u, v suitably.

3.1 Cauchy generators of the inverse

Let $(G, H) \in \mathbb{K}^{m \times \alpha} \times \mathbb{K}^{n \times \alpha}$ be $\nabla_{u,v}$ -generators of a matrix A , with $u = (u_1, \dots, u_m)$ and $v = (v_1, \dots, v_n)$. Let further r be the rank of A . Our goal is to decide if A has generic rank profile, and if so, to return generators $(Y, Z) \in \mathbb{K}^{r \times \alpha} \times \mathbb{K}^{r \times \alpha}$ of the inverse of the leading principal minor of A . Below, we write $\mu = \min(m, n)$.

For $0 \leq i \leq \mu$, write $A = \begin{bmatrix} A_{0,0}^{(i)} & A_{0,1}^{(i)} \\ A_{1,0}^{(i)} & A_{1,1}^{(i)} \end{bmatrix}$ with $A_{0,0}^{(i)} \in \mathbb{K}^{i \times i}$. If $A_{0,0}^{(i)}$ is invertible, set

$$S^{(i)} = \begin{bmatrix} S_{0,0}^{(i)} & S_{0,1}^{(i)} \\ S_{1,0}^{(i)} & S_{1,1}^{(i)} \end{bmatrix} = \begin{bmatrix} A_{1,1}^{(i)} - A_{1,0}^{(i)} A_{0,0}^{(i)-1} A_{0,1}^{(i)} & A_{1,0}^{(i)} A_{0,0}^{(i)-1} \\ -A_{0,0}^{(i)-1} A_{0,1}^{(i)} & A_{0,0}^{(i)-1} \end{bmatrix}$$

(this is similar to [14], where the order of block rows and columns was however different). The sequence of matrices $(S^{(i)})_{i=0 \dots \mu}$ starts from $A = S^{(0)}$ and ends at $A^{-1} = S^{(\mu)}$, at least when A is square and invertible. To understand where these matrices come from, we introduce the matrix $R^{(0)} = \begin{bmatrix} A \\ I_{\mu,n} \end{bmatrix}$ in $\mathbb{K}^{(m+\mu) \times n}$, where $I_{\mu,n} \in \mathbb{K}^{\mu \times n}$ is the rectangular matrix with ones on the main diagonal. If one applies to $R^{(0)}$ the column operations that reduce its first i th rows $[A_{0,0}^{(i)} \ A_{0,1}^{(i)}]$ to $[0 \ I_{i,i}]$ using $A_{0,0}^{(i)}$ as pivot, we get

$$R^{(i)} = \begin{bmatrix} S^{(i)} \\ I_{\mu,n} \end{bmatrix} = P^i R^{(0)} \begin{bmatrix} -A_{0,0}^{(i)-1} A_{0,1}^{(i)} & A_{0,0}^{(i)-1} \\ I_{n-i, n-i} & 0 \end{bmatrix}, \quad (4)$$

P being a cyclic permutation matrix that moves the rows up once.

Given integers a, b , $u_{a:b}$ denotes the sequence $(u_a, \dots, u_b)$, and similarly for $v_{a:b}$; we then define $u^{(i)} = (u_{i+1:m}, v_{1:i})$ and $v^{(i)} = (v_{i+1:n}, u_{1:i})$. A key result for the sequel is Lemma 3.1, which is [14, Proposition 1]; it gives $\nabla_{u^{(i)}, v^{(i)}}$ -generators of $S^{(i)}$ (remark that this operator is invertible, in view of our assumption on u and v).

For i as above, decompose G, H as $G = \begin{bmatrix} G_0^{(i)} \\ G_1^{(i)} \end{bmatrix}$, $H = \begin{bmatrix} H_0^{(i)} \\ H_1^{(i)} \end{bmatrix}$ with

$G_0^{(i)}, H_0^{(i)} \in \mathbb{K}^{i \times \alpha}$ and define $(Y^{(0)}, Z^{(0)}) = (G, H)$ and

$$Y^{(i)} = \begin{bmatrix} -A_{1,0}^{(i)} A_{0,0}^{(i)-1} G_0^{(i)} + G_1^{(i)} \\ -A_{0,0}^{(i)-1} G_0^{(i)} \end{bmatrix}, Z^{(i)} = \begin{bmatrix} -A_{0,1}^{(i)} A_{0,0}^{(i)-1} H_0^{(i)} + H_1^{(i)} \\ A_{0,0}^{(i)-1} H_0^{(i)} \end{bmatrix}.$$

LEMMA 3.1. $(Y^{(i)}, Z^{(i)})$ are $\nabla_{u^{(i)}, v^{(i)}}$ -generators for $S^{(i)}$.

If A has generic rank profile, this shows that for $r = \text{rank}(A)$, $(Y_0^{(r)}, Z_0^{(r)})$ are $\nabla_{v', u'}$ -generators for $A_{0,0}^{(r)-1}$, for u', v' as in Problem B, so they solve our problem; it remains to explain how to compute these matrices. Let i, j be non-negative integers with $0 \leq i+j \leq \mu$, and $A_{0,0}^{(i)}$ invertible. Decompose $S^{(i)}, Y^{(i)}$ and $Z^{(i)}$ into blocks

$$S^{(i)} = \begin{bmatrix} S_{0,0}^{(i,j)} & S_{0,1}^{(i,j)} \\ S_{1,0}^{(i,j)} & S_{1,1}^{(i,j)} \end{bmatrix}, Y^{(i)} = \begin{bmatrix} Y_0^{(i,j)} \\ Y_1^{(i,j)} \end{bmatrix}, Z^{(i)} = \begin{bmatrix} Z_0^{(i,j)} \\ Z_1^{(i,j)} \end{bmatrix}$$

with $S_{0,0}^{(i,j)} \in \mathbb{K}^{j \times j}$ and $Y_0^{(i,j)}, Z_0^{(i,j)} \in \mathbb{K}^{j \times \alpha}$, write $u^{(i)} = [u_0^{(i,j)} \ u_1^{(i,j)}]$ with $u_0^{(i,j)} \in \mathbb{K}^j$ and the same for $v^{(i)}$. Operations that derive $S^{(i+j)}$ from $S^{(j)}$ and $S^{(i)}$ from $A = S^{(0)}$ are similar; a direct calculation shows

$$S^{(i+j)} = \begin{bmatrix} S_{1,1}^{(i,j)} - S_{1,0}^{(i,j)} S_{0,0}^{(i,j)-1} S_{0,1}^{(i,j)} & S_{1,0}^{(i,j)} S_{0,0}^{(i,j)-1} \\ -S_{0,0}^{(i,j)-1} S_{0,1}^{(i,j)} & S_{0,0}^{(i,j)-1} \end{bmatrix}. \quad (5)$$

This implies the following formulae on the generators:

$$\begin{aligned} Y_1^{(i+j, m-j)} &= -S_{0,0}^{(i,j)-1} Y_0^{(i,j)}, & Z_1^{(i+j, n-j)} &= S_{0,0}^{(i,j)-t} Z_0^{(i,j)}, \\ Y_0^{(i+j, m-j)} &= -S_{1,0}^{(i,j)} Y_1^{(i+j, m-j)} + Y_1^{(i,j)}, & \\ Z_0^{(i+j, n-j)} &= -S_{0,1}^{(i,j)} Z_1^{(i+j, n-j)} + Z_1^{(i,j)}, \end{aligned} \quad (6)$$

this generalizes previous formulae for $(Y^{(i)}, Z^{(i)})$. The idea behind Eq. (6) is that by decomposing the column operations that reduce $[A_{0,0}^{(i+j)} \ A_{0,1}^{(i+j)}]$ to $[0 \ I_{i+j, i+j}]$ (see Eq.(4)) into two operations that respectively reduce the first i th rows using $A_{0,0}^{(i)}$ as pivot, then the next j rows using $S_{0,0}^{(i,j)}$ as pivot, we split the computation of $R^{(i+j)}$ into two similar operations that map $R^{(0)}$ to $R^{(i)}$ and $R^{(i)}$ to $R^{(i+j)}$.

Finally, to solve Problem B, we need to detect when A has generic rank profile. By construction, $S_{0,0}^{(i,j)}$ is the Schur complement of $A_{0,0}^{(i)}$, seen as a submatrix of $A_{0,0}^{(i+j)}$. This implies the following result.

LEMMA 3.2. If $A_{0,0}^{(i)}$ is invertible and has generic rank profile, $A_{0,0}^{(i+j)}$ has generic rank profile iff $S_{0,0}^{(i,j)}$ does, and $\text{rank}(A_{0,0}^{(i+j)}) = i + \text{rank}(S_{0,0}^{(i,j)})$.

3.2 A faster iterative algorithm

Using the previous discussion, we describe in Algorithm 1 a quadratic-time iterative algorithm for Problem B; we will use a step size $\beta \in \{1, \dots, \alpha\}$, given as a parameter. The formulae in (6) that compute $(Y^{(i+j)}, Z^{(i+j)})$ from $(Y^{(i)}, Z^{(i)})$ only involve $S_{0,0}^{(i,j)-1}, S_{0,1}^{(i,j)}, S_{1,0}^{(i,j)}$. Hence, line 8 recovers $S_{0,1}^{(i,j)}$ from its $\nabla_{u_0^{(i,j)}, v_1^{(i,j)}}$ -generators $(Y_0^{(i,j)}, Z_1^{(i,j)})$ and does the same for $S_{1,0}^{(i,j)}$ (Lemma 3.1); this allows us to apply (6).

For the correction of the algorithm, note that A has generic rank profile iff the rank of A is $r_{\max} = \max\{r \mid \forall \ell \leq r, \det(A_{0,0}^{(\ell)}) \neq 0\}$.

Algorithm 1: Iterative algorithm lter for Problem B

Input : Generators (G, H) of A and step size β
Output : $r = \text{rank}(A)$, generators of $A_{0,0}^{(r)-1}$

- 1 $i = 0, Y^{(0)} = G, Z^{(0)} = H, \bar{r} = \mu$
- 2 **while** $i \neq \bar{r}$ **do**
- 3 $j = \min(\beta, \bar{r} - i)$
- 4 Recover $S_{0,0}^{(i,j)}$ from its generators $(Y_0^{(i,j)}, Z_0^{(i,j)})$
- 5 Compute the rank ρ , rank profile, inverse of $S_{0,0}^{(i,j)}$
- 6 **if** $S_{0,0}^{(i,j)}$ *does not have generic rank profile* **then raise error**
- 7 **if** $\rho < j$ **then** $j = \rho, \bar{r} = i + \rho$
- 8 Recover $S_{0,1}^{(i,j)}, S_{1,0}^{(i,j)}$ from their generators
- 9 Compute the generators $(Y^{(i+j)}, Z^{(i+j)})$ of $S^{(i+j)}$
- 10 $i = i + j$
- 11 Recover $S^{(\bar{r})}$ from its generators $(Y^{(\bar{r})}, Z^{(\bar{r})})$
- 12 Read in $S^{(\bar{r})}$ the Schur complement of $A_{0,0}^{(\bar{r})}$ in A
- 13 **if** the Schur complement is non zero **then raise error**
- 14 **else return** $(\bar{r}, Y_1^{(\bar{r}, m-\bar{r})}, Z_1^{(\bar{r}, n-\bar{r})})$

The while loop of the algorithm will find the maximal leading principal minor of A that is non-zero and has generic rank profile. If at some step i , we know that $A_{0,0}^{(i)}$ is invertible and has generic rank profile, we use Lemma 3.2 to test if $A_{0,0}^{(i+j)}$ has the same properties.

The loop exits on one of two conditions that both ensures that $\bar{r} = r_{\max}$. Either $\text{rank}(S_{0,0}^{(i,j)})$ is always j , or $\text{rank}(S_{0,0}^{(i,j)}) < j$ at some point and after setting $\bar{r} = i + \text{rank}(S_{0,0}^{(i,j)})$, we have $\det(A_{0,0}^{(\ell)}) \neq 0$ for all $\ell \leq \bar{r}$ and $\det(A_{0,0}^{(\bar{r}+1)}) = 0$; we exit the while loop right after. After the while loop, it remains to test if $\text{rank}(A) = \bar{r}$, which is done by checking that the Schur complement of $A_{0,0}^{(\bar{r})}$ in A is zero.

Computing the rank profile at line 5 costs $O(\beta^\omega)$. Using block matrix multiplication in Eq. (2), we recover the matrices of line 8 in time $O(\beta^{\omega-2}\alpha(m+n))$, since $j \leq \beta \leq \alpha$. To compute $Y^{(i+j)}$ at line 9, we first compute $S_{0,0}^{(i,j)-1}$ in time $O(\beta^\omega)$, $S_{0,0}^{(i,j)-1}Y_1^{(i,j)}$ in time $O(\beta^{\omega-1}\alpha)$; then all other operations take time $O(\beta^{\omega-2}\alpha m)$. Similarly, computing $Z^{(i+j)}$ takes time $O(\beta^{\omega-2}\alpha n)$. We recover $S^{(\bar{r})}$ at line 11 from its $\nabla_{u^{(\bar{r})}, v^{(\bar{r})}}$ -generators $(Y^{(\bar{r})}, Z^{(\bar{r})})$ by means of (2), with a cost $O(\alpha^{\omega-2}mn)$ using block matrix multiplication.

One iteration of the while loop costs $O(\beta^{\omega-2}\alpha(m+n))$; we iterate $O(\mu/\beta)$ times, for a total cost of $O(\beta^{\omega-3}\alpha\mu(m+n))$, which is $O(\beta^{\omega-3}\alpha mn)$. This dominates the cost of line 11, so that the whole running time is $O(\beta^{\omega-3}\alpha mn)$. The algorithm of [35] uses $\beta = 1$, for which the cost is $O(\alpha mn)$; choosing $\beta = \alpha$, we benefit from fast matrix multiplication, as the cost drops to $O(\alpha^{\omega-2}mn)$.

3.3 The divide-and-conquer algorithm

We finally review the divide-and-conquer approach to solving Problem B, as a basis for the discussion of the next subsection. Algorithm 2 follows [26], recast in our framework, with the minor difference that do not assume A invertible, and that we explicitly check if A satisfies the generic rank profile assumption.

Recursive calls on lines 3, 8 will raise an error if the respective input matrix $A_{0,0}^{(i)}$ and its Schur complement $A_{1,1}^{(i)} - A_{1,0}^{(i)}A_{0,0}^{(i)-1}A_{0,1}^{(i)}$ are not in generic position, so we assume that we are not in that

Algorithm 2: Divide-And-Conquer algorithm DAC

Input : Generators (G, H) of A , threshold v_0
Output : $r = \text{rank}(A)$, generators of $A_{0,0}^{(r)-1}$

- 1 $\mu = \min(m, n), i = \lceil \mu/2 \rceil, (Y^{(0)}, Z^{(0)}) = (G, H)$
- 2 **if** $\max(m, n) \leq v_0$ **then return** lter(G, H, α)
- 3 $(r_0, Y_1^{(r_0, m-r_0)}, Z_1^{(r_0, n-r_0)}) = \text{DAC}(Y_0^{(0,i)}, Z_0^{(0,i)}, v_0)$
- 4 Recover Schur compl. gen. $(Y_0^{(r_0, m-r_0)}, Z_0^{(r_0, n-r_0)}) = (Y', Z')$ from $(Y_1^{(r_0, m-r_0)}, Z_1^{(r_0, n-r_0)})$ and $(Y^{(0)}, Z^{(0)})$
- 5 **if** $r_0 < i$ **then**
- 6 **if** Schur complement is non zero **then raise error**
- 7 **else return** $(r_0, Y_0^{(r_0)}, Z_0^{(r_0)})$
- 8 $(r_1, Y_1^{(r_1, m-r_1)}, Z_1^{(r_1, n-r_1)}) = \text{DAC}(Y', Z', v_0), r_2 = r_0 + r_1$
- 9 **return** r_2 and $(Y_1^{(r_2, m-r_2)}, Z_1^{(r_2, n-r_2)})$ computed from $(Y_1^{(r_2, m-r_1)}, Z_1^{(r_2, n-r_1)})$ and $(Y^{(r_0)}, Z^{(r_0)})$

case; then, the correction proof is similar to that in Section 3.2. Computations of lines 4, 9 are of the same essence; using Eq. (5) and (6), we recover the full generators $(Y^{(i+j)}, Z^{(i+j)})$ of $S^{(i+j)}$ if we know the generators $(Y_1^{(i+j, m-j)}, Z_1^{(i+j, n-j)})$ of $S_{0,0}^{(i,j)-1}$ and previous full generators $(Y^{(i)}, Z^{(i)})$ of $S^{(i)}$. Rather than reconstructing the submatrices of $S^{(i,j)}$ as in the previous subsection, we use them through their generators: the formulae of (6) boil down to $O(1)$ multiplications of Cauchy-like matrices (for which we have generators of length α) by α vectors.

In general, these multiplications cost $O(\alpha^2 \mathcal{M}(v) \log(v))$ operations, with $v = \max(m, n)$. As pointed out in [39, Theorem 5.3.1] (see also [15]), if u and v are geometric progressions, the cost for the matrix-vectors multiplication drops to $O(\alpha^2 \mathcal{M}(v))$. If u and v have the same ratio (if this is the case at the top-level, this will remain the case for all recursive calls), by Remark 2.1, we can save a further constant factor. For large values of α , we can also apply the algorithm of [8], which uses fast polynomial matrix multiplication to reduce the cost to $O(\alpha^{\omega-1} \mathcal{M}(v))$.

On line 6, the Schur complement is zero iff $Y'Z' = 0$. This is tested by finding a minimal set of independent rows in Z' (this takes time $O(\alpha^{\omega-1}v)$) and multiplying their transposes by Y' . There are at most α such rows, so this takes time $O(\alpha^{\omega-1}v)$ as well. Taking all recursive steps into account, the total cost of DAC is a factor $\log(v)$ times that of the Cauchy-like matrix products.

3.4 Experimental results

We implemented the algorithms described so far, together with all subroutines they rely on, in a C++ library available at <https://seafire.lirmm.fr/f/01b53dc420/>. Our implementation is based on Shoup's NTL [43, 44] version 10.3.0, and is dedicated to word-size primes (NTL's lzz_p class); the divide-and-conquer algorithm of Subsection 3.3 actually requires FFT primes. In all that follows, timings are measured on an Intel i7-4790 CPU with 32 GB RAM; only one thread is used throughout.

Some goals of our experiments are to assess whether fast matrix multiplication can bring practical improvements that reflect the theoretical ones, what are reasonable crossover points between iterative and divide-and-conquer algorithms, and what is the range of applicability of these structured methods compared to dense

linear algebra. Thus, we focused on comparisons between our own implementations of the various techniques seen so far, and did not attempt comparisons with other systems.

The main subroutines we need are polynomial and matrix computations. NTL already offers efficient FFT-based polynomial arithmetic; matrix multiplication over small fields $\mathbb{Z}/p\mathbb{Z}$, for $p < 2^{23}$, is now extremely efficient, comparable to reference implementations such as FFLAS-FFPACK [19]. On top of this, we implemented a fast polynomial matrix multiplication, using the cyclotomic TFT of [1].

We first discuss Problem B. Our first tests compare our new $O(\alpha^{\omega-2}n^2)$ algorithm to that of [35], with runtime $O(\alpha n^2)$. Except for very small values of n , say $n < 50$, for which the behavior fluctuates rapidly, we found that the new algorithm becomes more efficient for rather small values of α : our crossover points are $\alpha = 8$ or 9 for primes less than 2^{23} , and $\alpha = 13$ or 14 for larger word-size primes. The following graph shows the time ratio between the algorithm of [35] and our algorithm, for $\alpha = 30$; the larger value of p used here is the FFT prime $p = 82705526964617217 = 7^2 2^{54} + 1$. We see that for small primes, for which matrix multiplication is very efficient, our algorithm brings a substantial improvement.

We consider next the divide-and-conquer algorithm. The key factor for the efficiency of this algorithm is the cost of multiplying an $n \times n$ Cauchy-like matrix of displacement rank α by α vectors. We compare the approach of cost $O(\alpha^2 \mathcal{M}(n))$ of Remark 2.1 to the algorithm of [8], with cost $O(\alpha^{\omega-1} \mathcal{M}(n))$, with a view of determining for what values of α (if any) the latter becomes useful.

For the former algorithm, because we are able to cache several FFTs, we found it slightly more advantageous to use NTL's FFT rather than TFTs. The runtime of the first algorithm then displays the typical FFT staircase behavior, so that as n grows, the crossover value for α fluctuates, roughly between 30 and 55. The following graph shows the time ratio between the algorithm of [8] and the direct one, for $\alpha = 60$; at best, the new algorithm wins by a factor of 2. The results do not depend much on the nature of the prime (since polynomial arithmetic is a significant part of the runtime, and behaves essentially in the same manner, independently of p).

Using these results, we determined empirical crossover values $v_0(\alpha)$ to end the recursion in the divide-and-conquer algorithm and switch to the iterative algorithm. We expect $v_0(\alpha)$ to grow with α , since it behaves like the solution of $\alpha n^2 = \alpha^2 \mathcal{M}(n) \log(n)$ (assuming here we do not use fast linear algebra, for simplicity). The following table reports some values for v_0 (obtained by searching for v_0 in increments of 100). The threshold is higher for small primes, since such primes benefit more from our new iterative algorithm.

α	1	5	10	15	20	25	30
$p = 65537$	400	400	500	1000	2000	3300	4200
$p \approx 2^{59}$	200	400	400	700	1300	1600	2000

These values being set, we show runtimes for solving Problem B modulo $p = 65537$ and $p = 82705526964617217$, for increasing values of n , with $\alpha = 5$ and $\alpha = 50$; we also show the runtime of a dense matrix inversion in the same size. For a small displacement rank such as $\alpha = 5$, the runtime is essentially the same for these two primes; with $\alpha = 50$, we observe a difference, by a factor of up to 3. In any case, there is a clear gain over dense methods.

We also determined, for a given value of n , the crossover value α_0 above which dense linear algebra becomes faster than structured methods. The value $\alpha_0(n)$ grows quite regularly with n , good approximations being $\alpha_0(n) \simeq 0.2n$ (for $p < 2^{23}$) and $\alpha_0(n) \simeq 0.25n$ (for larger values of p). This means that there is a wide range of inputs for which structured methods can be of use.

Our solution of Problem A is a direct reduction to the Cauchy-like case. Putting the problem into Cauchy form by means of the formulae of Eq. (3) accounts for a small fraction of the total runtime: between 5% and 10% for small α (say $\alpha < 10$), and less than 5% for larger values of α , in all instances we considered.

In the thousands of experiments we made, for matrix sizes such as $n \geq 5000$ and small primes such as $p = 65537$, we observed some instances where the Cauchy matrix did not have generic rank profile. This never happened for p having more than 50 bits.

4 MODULAR TECHNIQUES

We now address Problem A in the particular case where $\mathbb{K} = \mathbb{Q}$. Several *modular algorithms* are available to solve dense linear systems over $\mathbb{Q}$. A first option relies on the Chinese Remainder Theorem, solving the system modulo several primes $p_1, p_2, \dots$ before reconstructing the solution, making sure to ensure consistency of the modular solutions when $\ker(T)$ has dimension more than 1. Other approaches such as Dixon's algorithm [17], Newton iteration or divide-and-conquer algorithms use one prime p , and lift the solution modulo powers of p .

Newton iteration for structured matrices goes back to Pan's article [38] (for Toeplitz-like matrices), and is explained in detail in [39, Chapter 7]. To the best of our knowledge, the other approaches mentioned above have not been discussed in the literature on structured matrices, save for the second author's PhD thesis [32]. Our goal in this section is to first briefly present some of these techniques and analyze their complexity for the problem at hand; we will then discuss their practical performance. We will highlight in particular the case of algebraic approximants, for which we are able to obtain significant improvements.

We denote by $\mathcal{J} : \mathbb{N} \rightarrow \mathbb{R}$ a function such that integers of bit size at most d can be multiplied in $\mathcal{J}(d)$ bit operations; we can take $\mathcal{J}(d) = O(d \log(d) \log \log(d))$ or $\mathcal{J}(d) = d \log(d) 2^{O(\log^*(d))}$ [20, 41]. As in [21], we assume that $d \mapsto \mathcal{J}(d)/d$ is non-decreasing.

4.1 Solving square systems by lifting

In this subsection, we are given a prime power p^t , where t is a power of 2, an $n \times n$ matrix A and a vector b , both with entries modulo p^t , and we assume that A is invertible modulo p^t , with inverse B . We discuss algorithms that solve the equation $Ax = b$, and we estimate their complexity when A is a structured matrix. To simplify the cost analysis, we assume here that $p = O(1)$.

We take $\mathcal{C}_A$ such that, for any t , we can compute $Ax \bmod p^t$ in $O(\mathcal{C}_A \mathcal{J}(t))$ bit operations, given a vector x with entries defined modulo p^t . Below, we will assume that A is Cauchy-like as in Remark 2.1, and given by generators of length α , so that we can take $\mathcal{C}_A \in O(\alpha \mathcal{M}(n))$. We will see however that in the case of algebraic approximants, better estimates are available.

We consider two approaches: a divide-and-conquer algorithm and Newton iteration, which both feature a running time linear in the target precision t . We start with the divide-and-conquer approach. A version of it is in [2] for dense matrices; the PhD thesis [32] describes this algorithm for Toeplitz-like matrices.

Assume that we have obtained generators of length α for B , for instance by applying the algorithm of the previous section. Thus, at the leaves of the recursion, each product $Bb \bmod p$ can be computed using $O(\alpha \mathcal{M}(n))$ bit operations; using our assumption on $\mathcal{J}$, the total runtime is then $O(\mathcal{C}_A \mathcal{J}(t) \log(t) + \alpha \mathcal{M}(n)t)$ bit operations. With our upper bounds on $\mathcal{C}_A$, this simplifies further as $O(\alpha \mathcal{M}(n) \mathcal{J}(t) \log(t))$.

We turn next to Newton iteration. The matrix form of Newton iteration computes the inverses $B_k = A^{-1} \bmod p^{2^k}$ by $B_{k+1} =$

Algorithm 3: Divide-And-Conquer algorithm $\text{DAC}_{\mathbb{Q}}$

Input : $A, b, B \bmod p, p, t$ as above
Output : a solution of $Ax = b \bmod p^t$
1 **if** $t = 1$ **then return** $Bb \bmod p$
2 Compute $x_0 = \text{DAC}_{\mathbb{Q}}(A, b, B, p, t/2)$
3 Compute $r_0 = (Ax_0 - b) \bmod p^t$ and $r_1 = r_0/p^{t/2}$
4 Compute $x_1 = \text{DAC}_{\mathbb{Q}}(A, r_1, B, p, t/2)$
5 **return** $x_0 - p^{t/2}x_1 \bmod p^t$

$2B_k - B_kAB_k \bmod p^{2^{k+1}}$; once they are known, we deduce the solution to our system by means of a matrix-vector product.

Let (G, H) be generators for A . Pan's insight was to use the relation above to compute the generators $(X, Y) = (-BG, B^tH)$ for B . Given $(X_k, Y_k) = (X, Y) \bmod p^{2^k}$, we can reconstruct B_k , but to deduce (X_{k+1}, Y_{k+1}) , we have to multiply G and H by B_{k+1} or B_{k+1}^t . This is done using the expression for B_{k+1} above, which gives

$$\begin{aligned} X_{k+1} &= -(2B_k - B_kAB_k) G \bmod p^{2^{k+1}} \\ Y_{k+1} &= (2B_k - B_kAB_k)^t H \bmod p^{2^{k+1}}. \end{aligned}$$

This time, we are multiplying A and B_k (and their transposes) by matrices of size $n \times \alpha$, all computations being done modulo $p^{2^{k+1}}$. Each multiplication by *one* vector takes $O(\alpha \mathcal{M}(n) \mathcal{J}(2^k))$ operations (even if $\mathcal{C}_A$ is less than $O(\alpha \mathcal{M}(n))$, multiplications by B_k are the bottleneck). Since we multiply these matrices by α vectors, taking all steps into account, we arrive at a total of $O(\alpha^2 \mathcal{M}(n) \mathcal{J}(t))$ bit operations. Using the algorithm of [8], this can be further reduced to $O(\alpha^{\omega-1} \mathcal{M}(n) \mathcal{J}(t))$; this improvement was not implemented.

Altogether, because it computes (generators of) a whole inverse, Newton iteration is slower by a factor α (or slightly less, if we use [8]); on the other hand, it saves a factor of $\log(t)$. This is similar to what one observed when comparing these techniques for e.g. the solution of differential equations [5, 7].

4.2 Solving Problem A

The techniques above allow us to solve instances of Problem A, that is, find nullspace elements for a mosaic Toeplitz matrix T , but they would require the knowledge of a maximal invertible minor of it. However, there is no guarantee that T possesses such a minor that would be mosaic Toeplitz as well, and we do not know how to find it efficiently. Hence, we rely on the transformation to the Cauchy structure / regularization technique of Section 2.

Over $\mathbb{Q}$, this approach has a certain shortcoming: the output is a vector $b = W_v \begin{bmatrix} -A_{0,0}^{(r)-1} A_{0,1}^{(r)} c \\ c \end{bmatrix}$, with $A = V_u T W_v$, $A_{0,0}^{(r)}$ a maximal minor of it and c a random vector in $\mathbb{K}^{n-r}$, and due to the preconditioning, the entries of b are expected to be of large height (larger than what we may expect for a solution of T). When $\ker(T)$ has dimension 1, we are not affected by this issue. Indeed, in this case, all solutions are of the form λb_0 , for some vector $b_0 \in \mathbb{Z}^n$ whose bit-size can be bounded only in terms of T . Hence, it suffices to compute the solution b of the regularized system modulo a large enough integer N and normalize it by setting one of its entries to 1; this gives us the normalization of $b_0 \bmod N$.

This idea could be extended to the case of an arbitrary nullspace dimension, but it would be tantamount to lifting a whole nullspace basis. Instead, we will reduce the nullspace dimension by adding

a new block of equations; when $\ker(T)$ has moderate dimension, this barely affects the overall runtime (for particular applications to algebraic approximants, another solution is described below).

For simplicity, we discuss here a version of the algorithm that may run forever in unlucky cases. We choose a prime p and compute $\nabla_{u,v}$ -generators for $A = V_u T W_v \pmod{p}$, for u, v as in Remark 2.1. We call the algorithm of the previous section, in order to determine the rank of A ; if A does not have generic rank profile, choose another u, v . After an expected $O(1)$ attempts, we obtain the rank of A (as a matrix over $\mathbb{F}_p$), and thus the dimension s of its nullspace. If s is greater than one, we add a block of Toeplitz matrices having $s - 1$ rows to T , with small random entries, and update A accordingly. Heuristically the new matrix T has nullspace dimension 1 (otherwise, add another block of equations).

Let $d = \text{rank}(T)$, A_d be the $d \times d$ top-left submatrix of A , and b_d be the vector of the first d entries of the last column of A . We assume that A has generic rank profile, so that A_d is invertible. We compute $x = A_d^{-1} b_d \pmod{p}$ and $y = W_v[x_1, \dots, x_d, -1]^t \pmod{p}$, we normalize y by dividing it by its first non-zero entry, and we set $t = 1$. While we either cannot apply rational reconstruction to the entries of y , or after applying rational reconstruction to y we have $Ty \neq 0$, we do the following: set $t = 2t$, update modulo p^t the quantities A_d, b_d, x (use one of the algorithms of Section 4.1) and y , and divide y by its first non-zero entry.

A complete analysis of this algorithm would quantify the primes of bad reduction, give bounds that allow us to stop lifting the solutions if we reduced modulo such a bad prime and study the previous reduction to nullity one of T ; we leave this to future work. In any case, if the lifting stops, we have obtained a solution to our system. Due to the doubling nature of this procedure, the runtime is proportional to that of the algorithm for solving square systems used at line 3, plus the cost of rational reconstruction. The former is $O(\mathcal{C}_A \mathcal{J}(t) \log(t) + \alpha \mathcal{M}(n)t)$ bit operations using divide-and-conquer, and $O(\alpha^2 \mathcal{M}(n) \mathcal{J}(t))$ or $O(\alpha^{\omega-1} \mathcal{M}(n) \mathcal{J}(t))$ using Newton iteration; the latter is $O(n \mathcal{J}(t) \log(t))$.

Finally, we mention the cost of Chinese Remaindering techniques: instead of computing the solution modulo the t -th power of a single prime p , we might want to solve the system modulo t primes of the same magnitude. If we assume that A remains invertible with generic rank profile modulo all these primes, and all these primes are $O(1)$, the runtime for solving the systems is now $O(\alpha^2 \mathcal{M}(n) \log(n)t)$ or $O(\alpha^{\omega-1} \mathcal{M}(n) \log(n)t)$, by the results of Section 3 (to this, we have to add the cost $O(n \mathcal{J}(t) \log(t))$ of Chinese Remaindering and rational reconstruction).

We conclude this subsection with an important particular case, the computation of algebraic approximants. We are given a power series f in $\mathbb{Q}[[x]]$, together with degree bounds d, e ; our goal is to compute a polynomial $P \in \mathbb{Q}[x, y]$, with $\deg(P, x) \leq d, \deg(P, y) \leq e$, such that $P(x, f) = 0$. For any $\sigma \geq 0$, finding P with the above degree bounds and such that $P(x, f) = 0 \pmod{x^\sigma}$ is a Hermite-Padé approximation problem of a very special kind (all input series are powers of f). To our knowledge, no algorithm in the framework of Section 3 can exploit this extra structure; in the case of computations over $\mathbb{Q}$, we will now see that improvements are possible.

First, we show how to simplify the reduction to nullity one. Théorème 7.15 in [6] shows that if such a P exists and is irreducible, then any $Q \in \mathbb{Q}[x, y]$ with the same degree bounds as above and

such that $Q(x, f) = 0 \pmod{x^{2de+1}}$ is a multiple of P . From this, one deduces easily that if we compute two such polynomials Q_1, Q_2 , their GCD will generically be P itself. For all primes p except a finite number, the rank of our matrix T does not change modulo p , and the GCD of two random basis elements commutes with reduction modulo p (note that $P \pmod{p}$ may not be irreducible anymore). Hence, we can (probabilistically) find $P \pmod{p}$ by computing two solutions Q_1, Q_2 to the above Hermite-Padé problem modulo p and taking their GCD. This reveals the support of P ; we can then refine the degree bounds in our Hermite-Padé problem, which in turn reduces the nullity of matrix T to 1.

Next, we show how to speed-up algorithm DAC_Q in this case. The block-Toeplitz matrix T of our Hermite-Padé problem has displacement rank $\alpha = O(e)$, with $e + 1$ Sylvester blocks having $O(d)$ columns and $O(de)$ rows; as a result, the naive estimate on the cost of the matrix-vector product by matrix $A = V_u T W_v$ is $\mathcal{C}_A = O(e \mathcal{M}(de))$. However, we can reduce this cost using baby-steps / giant steps techniques: the bivariate modular composition algorithm of [36] shows that we can do matrix-vector product by T using $O(e^{1/2} \mathcal{M}(de) + e^{(\omega+1)/2} \mathcal{M}(d))$ operations; since multiplications by V_u and W_v take $O(\mathcal{M}(de))$, we obtain the improved estimate $\mathcal{C}_A = O(e^{1/2} \mathcal{M}(de) + e^{(\omega+1)/2} \mathcal{M}(d))$ in this case. Algorithm DAC_Q is the only algorithm we know of that takes into account the extra structure of algebraic approximants; we expect that similar improvements are possible for *differential approximants*, using the evaluation algorithm of [12].

4.3 Experimental Results

Our experiments are dedicated to solving instances of Hermite-Padé approximation, which is a useful particular case of Problem A. We discuss two families of problems, first the approximation of general power series, then algebraic approximants. Timings are measured on the same machine as Subsection 3.4.

In both cases, we show two graphs: on the left, we have five blocks, each with n columns; on the right, we have n blocks, each with n columns. With the notation of the introduction, this means we are looking for approximants $(p_0, \dots, p_4)$, resp. $(p_0, \dots, p_{n-1})$, with degree bounds (n, n, n, n, n) , resp. $(n, \dots, n)$. The displacement rank α of these matrices is 6, resp. $n + 1$.

We first examine the results for general power series. Our experiments showed that the runtime grows predictably with respect to the input bit-size; as a result, we fix the input coefficients to be 10 bit integers, so the number of lifting steps depends just on n .

Matrices are generated so as to have 1 less row than they have columns; since the inputs are random, they have nullspace of dimension 1 (we also generated instances with nullity up to 10; using our heuristic to reduce nullity, runtimes were almost indistinguishable). The primes we use have 59 to 60 bits. For DAC_Q and Newton iteration, the sharp increases indicate an additional lifting step.

Newton iteration is slower than DAC_Q , especially when the number of block grows (since its runtime is quadratic in the displacement rank α). Newton iteration should theoretically be competitive with DAC_Q when the number of blocks is fixed, and the size (and thus the output bit size) grow, since its theoretical runtime is better by a $\log(t)$ factor, where t is essentially the output bit size. However, this is not noticeable on our experiments: in practice, the integer multiplication function $\mathcal{I}(d)$ grows like $d^{1+\varepsilon}$, for some $\varepsilon > 0$; in that case, the analysis of DAC_Q can be refined to $O(\alpha \mathcal{M}(n) \mathcal{I}(t))$.

CRT, on the other hand, seems to be competitive with DAC_Q when α is small but is significantly worse as α grows: the runtime CRT is not linear in α , while DAC_Q is.

Next, we examine the computation of algebraic approximants; on the basis of the previous experiments, we consider algorithm DAC_Q only. We start by generating a bivariate polynomial $P(x, y)$ and compute one of its power series solutions f . Since we expect the coefficients of P to be smaller than the coefficients of f , we can better control the behavior of the algorithms by choosing the bit size of P (here, it was fixed to be 1000 bit integers). We compare algorithm DAC_Q as in the general case, with $\mathcal{C}_A = O(\alpha \mathcal{M}(n))$, to the improved version using the bivariate modular composition (BMC) algorithm of [36] described above, featuring a lower value for $\mathcal{C}_A$. The latter algorithm uses polynomial matrix multiplication, which we implemented using reduction modulo FFT primes and TFT polynomial multiplication.

As a result, we can see a significant difference between the two algorithms as the size of the matrix grows: the theoretical speed-up predicted in Subsection 4.2 is observed in practice.

Overall, the divide-and-conquer lifting algorithm turned out to be the most efficient method in all our experiments, especially as it can take extra structure into account, as in the case of algebraic approximants.

REFERENCES

- [1] A. Arnold and É. Schost. A Truncated Fourier Transform middle product. *ACM Commun. Comput. Algebra*, 48(3/4):98–99, 2015.
- [2] J. Berthomieu and R. Lebreton. Relaxed p -adic Hensel lifting for algebraic systems. In *ISSAC'12*, pages 59–66. ACM, 2012.
- [3] R. R. Bitmead and B. D. O. Anderson. Asymptotically fast solution of Toeplitz and related systems of linear equations. *Linear Algebra Appl.*, 34:103–116, 1980.
- [4] L. I. Bluestein. A linear filtering approach to the computation of the discrete Fourier transform. *IEEE Trans. Electroacoustics*, AU-18:451–455, 1970.
- [5] A. Bostan, M. F. I. Chowdhury, R. Lebreton, B. Salvy, and É. Schost. Power series solutions of singular (q)-differential equations. In *ISSAC'12*, pages 107–114. ACM, 2012.
- [6] A. Bostan, F. Chyzak, M. Giusti, R. Lebreton, G. Lecerf, B. Salvy, and É. Schost. Algorithmes efficaces en calcul formel. hal.archives-ouvertes.fr/AECF/, 2017.
- [7] A. Bostan, F. Chyzak, F. Ollivier, B. Salvy, S. Sedoglavic, and É. Schost. Fast computation of power series solutions of systems of differential equations. In *Symposium on Discrete Algorithms, SODA'07*, pages 1012–1021. ACM-SIAM, 2007.
- [8] A. Bostan, C.-P. Jeannerod, C. Moulleron, and É. Schost. On matrices with displacement structure: generalized operators and faster algorithms. preprint.
- [9] A. Bostan, C.-P. Jeannerod, and É. Schost. Solving structured linear systems with large displacement rank. *Theor. Comput. Sci.*, 407(1):155–181, 2008.
- [10] A. Bostan, G. Lecerf, and É. Schost. Tellegen's principle into practice. In *ISSAC'03*, pages 37–44. ACM, 2003.
- [11] A. Bostan and É. Schost. Polynomial evaluation and interpolation on special sets of points. *J. Complexity*, 21(4):420–446, 2005.
- [12] A. Bostan and É. Schost. Fast algorithms for differential equations in positive characteristic. In *ISSAC'09*, pages 47–54. ACM, 2009.
- [13] D. G. Cantor and E. Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28(7):693–701, 1991.
- [14] J.-P. Cardinal. On a property of Cauchy-like matrices. *C. R. Acad. Sci. Paris Série I*, 388:1089–1093, 1999.
- [15] Z. Chen and V. Y. Pan. An efficient solution for Cauchy-like systems of equations. *Computers Math. Applic.*, 48:529–537, 2004.
- [16] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251–280, March 1990.
- [17] J. Dixon. Exact solution of linear equations using p -adic expansions. *Numer. Math.*, 40:137–141, 1982.
- [18] J. Durbin. The fitting of time series models. *Rev. Inst. Int. Stat.*, 28:233–243, 1960.
- [19] FFLAS-FPACK-Team. *FFLAS-FPACK: Finite Field Linear Algebra Subroutines / Package*, v2.2.2 edition, 2016. <http://github.com/linbox-team/fflas-ffpack>.
- [20] M. Fürer. Faster Integer Multiplication. In *STOC'07*, pages 57–66, 2007.
- [21] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, Cambridge, third edition, 2013.
- [22] G. Hanrot, M. Quercia, and P. Zimmermann. The Middle Product Algorithm. *I. Appl. Algebra Engrg. Comm. Comp.*, 14(6):415–438, 2004.
- [23] G. Heinig and T. Amdeberhan. On the inverses of Hankel and Toeplitz mosaic matrices. In *Seminar Analysis (Berlin, 1987/1988)*, pages 53–65, 1988.
- [24] T. Huckle. Implementation of a superfast algorithm for symmetric positive definite linear equations of displacement rank 2. volume 2296 of *Proceedings of the SPIE*, pages 494–503, 1994.
- [25] O. H. Ibarra, S. Moran, and R. Hui. A generalization of the fast LUP matrix decomposition algorithm and applications. *J. Algorithms*, 3(1):45–56, 1982.
- [26] C.-P. Jeannerod and C. Moulleron. Computing specified generators of structured matrix inverses. In *ISSAC'10*, pages 281–288. ACM, 2010.
- [27] T. Kailath, I. Gohberg, and V. Olshevsky. Fast Gaussian elimination with partial pivoting for matrices with displacement structure. *Math. Comp.*, 64(212):1557–1576, 1995.
- [28] T. Kailath, S. Y. Kung, and M. Morf. Displacement ranks of matrices and linear equations. *J. Math. Anal. Appl.*, 68(2):395–407, 1979.
- [29] T. Kailath and A. H. Sayed. *Fast Reliable Algorithms for Matrices with Structure*. SIAM, 1999.
- [30] E. Kaltofen. Asymptotically fast solution of Toeplitz-like singular linear systems. In *ISSAC'94*, pages 297–304. ACM, 1994.
- [31] F. Le Gall. Powers of tensors and fast matrix multiplication. In *ISSAC'14*, pages 296–303. ACM, 2014.
- [32] R. Lebreton. *Contributions to relaxed algorithms and polynomial system solving*. PhD thesis, École Polytechnique, Palaiseau, France, 2012.
- [33] N. Levinson. The Wiener RMS error criterion in filter design and prediction. *J. Math. Phys.*, 25:261–278, 1947.
- [34] M. Morf. Doubling algorithms for toeplitz and related equations. In *IEEE Conference on Acoustics, Speech, and Signal Processing*, pages 954–959, 1980.
- [35] C. Moulleron. Algorithmes rapides pour la résolution de problèmes algébriques structurés. Master's thesis, ENS Lyon, 2008.
- [36] M. Nüsken and M. Ziegler. Fast multipoint evaluation of bivariate polynomials. In *ESA 2004*, number 3222 in LNCS, pages 544–555. Springer, 2004.
- [37] V. Y. Pan. On computations with dense structured matrices. *Math. Comp.*, 55(191):179–190, 1990.
- [38] V. Y. Pan. Parametrization of Newton's iteration for computations with structured matrices and applications. *Computers Math. Applic.*, 24(3):61–75, 1992.
- [39] V. Y. Pan. *Structured Matrices and Polynomials*. Birkhäuser Boston Inc., 2001.
- [40] V. Y. Pan and A. Zheng. Superfast algorithms for Cauchy-like matrix computations and extensions. *Linear Algebra Appl.*, 310:83–108, 2000.

- [41] A. Schönhage and V. Strassen. Schnelle Multiplikation großer Zahlen. *Computing*, 7:281–292, 1971.
- [42] H. Sexton, M. Shensa, and J. Speiser. Remark on a displacement-rank inversion method for Toeplitz systems. *Linear Algebra Appl.*, 45:127–130, 1982.
- [43] V. Shoup. NTL: A library for doing number theory. <http://www.shoup.net>.
- [44] V. Shoup. A new polynomial factorization algorithm and its implementation. *J. Symb. Comp.*, 20(4):363–397, 1995.
- [45] W. F. Trench. An algorithm for the inversion of finite Toeplitz matrices. *J. Soc. Indust. Appl. Math.*, 12:515–522, 1964.