



HAL
open science

In-place fast polynomial modular remainder

Jean-Guillaume Dumas, Bruno Grenet

► **To cite this version:**

Jean-Guillaume Dumas, Bruno Grenet. In-place fast polynomial modular remainder. Univ. Grenoble Alpes. 2023. hal-03979016v2

HAL Id: hal-03979016

<https://hal.science/hal-03979016v2>

Submitted on 24 Feb 2023 (v2), last revised 12 Jan 2024 (v6)

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.

In-place fast polynomial modular remainder

Jean-Guillaume Dumas* Bruno Grenet*

February 24, 2023

Abstract

We consider the computation of the euclidean polynomial modular remainder $R(X) \equiv A(X) \pmod{B(X)}$ with A and B of respective degrees n and $m \leq n$. If the multiplication of two polynomials of degree k can be performed with $\mathfrak{M}(k)$ operations and $\mathcal{O}(k)$ extra space, then standard algorithms for the remainder require $\mathcal{O}\left(\frac{n}{m}\mathfrak{M}(m)\right)$ arithmetic operations and, apart from that of A and B , $\mathcal{O}(n)$ extra memory. This extra space is notably usually used to store the whole quotient $Q(X)$ such that $A = BQ + R$ with $\deg R < \deg B$.

We avoid the storage of the whole of this quotient, and propose an algorithm still using $\mathcal{O}\left(\frac{n}{m}\mathfrak{M}(m)\right)$ arithmetic operations but only $\mathcal{O}(m)$ extra space.

When the divisor B is sparse with a constant number of non-zero terms, the arithmetic complexity bound reduces to $\mathcal{O}(n)$.

When it is allowed to use the input space of A or B for intermediate computations, but putting A and B back to their initial states after the completion of the remainder computation, we further propose an in-place algorithm (that is with its extra required space reduced to $\mathcal{O}(1)$ only) using $\mathcal{O}(n)$ arithmetic operations over any field of zero or odd characteristic and over most of the characteristic two ones.

To achieve this, we develop techniques for Toeplitz matrix operations which output is also part of the input.

1 Introduction

Modular methods with dense univariate polynomials over a finite ring are of central importance in computer algebra and symbolic computation. For instance, they are largely used with lifting or Chinese remaindering techniques as the bases to compute at a larger precision. There, the quotient of the Euclidean division is not needed, but is very often computed anyway along the algorithm.

In terms of arithmetic operations, from the work of [20, 16] to more recent results of [4], many sub-quadratic algorithms were developed for this task,

*Université Grenoble Alpes. Laboratoire Jean Kuntzmann, CNRS, UMR 5224. 700 avenue centrale, IMAG - CS 40700, 38058 Grenoble, cedex 9 France. {firstname.lastname}@univ-grenoble-alpes.fr

based on fast polynomial multiplication [6, 15, 25]. But these fast algorithms come at the expense of (potentially large) extra temporary space to perform the computation. On the contrary, classical, quadratic algorithms, when computed sequentially, quite often require very few (constant) extra registers. Further work then proposed simultaneously “fast” and “in-place” algorithms, for both matrix and polynomial operations [5, 11, 12].

We here propose algorithms to extend the latter line of work. In the polynomial setting, we compute the remainder only of the Euclidean division. This means that, e.g., with respect to [12, Alg. 3], we obtain the remainder without needing any space for the quotient. As polynomials and Toeplitz matrices are indeed different representations of the same objects, see, e.g., [2, 3, 11], we also develop methods for Toeplitz matrix operations, *in-place with accumulation* or *over-place*: those are difficult cases where the output is also part of the input. The difficulty when the result overwrites (parts) of the input, is that in-place methods start with absolutely no margin for extra space. Thus, for instance, the generic recursive techniques of [5, 11] usually do not apply.

Next we first detail our model for in-place computations in Section 1.1, recall some quadratic in-place algorithms in Section 1.2 and some fast algorithms in Section 1.3. Then, in Section 2, we derive new in-place algorithms for circulant and Toeplitz matrices over fields of zero or odd characteristic. In Section 3 we then propose a fast remaindering algorithm using only a multiple of the remainder space. Further, in Section 4, we extend this to using as extra space only a small fraction of the remainder space. Finally in Section 5 we combine these techniques to obtain fast in-place algorithms for the polynomial remainder only. We also propose in Appendix B some ways to adapt our algorithms to the characteristic two specificities.

1.1 In-place model

There exist different models for in-place algorithms. We here choose to call *in-place* an algorithm using only **the space of its inputs, its outputs, and $\mathcal{O}(1)$ extra space**. But algorithms are only allowed to modify their inputs, **provided that their inputs are restored to their initial state** afterwards. We do not take into account the potential space used by the recursive call stack. (This space is bounded by the recursion depth of the algorithms, and managed in practice by the compiler.) This is a less powerful model than when the input is purely read-only, but it turns out to be crucial in our case, especially when we have accumulation operations (e.g., like $\vec{a} += \vec{b}$). If atomic operations are used to perform the accumulation, then usually a constant number of extra space is sufficient. Differently, fast algorithms usually work with (non-constant) large blocks of data and having only the latitude to use a constant extra-space is often not sufficient.

Also, as a variant, by *over-place*, we mean an algorithm where the output replaces (parts) of its input (e.g., like $\vec{a} = b \cdot \vec{a}$). Similarly, we allow all of the input to be modified, provided that the parts of the input that are not the output are restored afterwards. In the following we signal by a “**Read-only:**”

tag the parts of the input that the algorithm is not allowed to modify (the other parts are modifiable as long as they are restored).

The main limitations of this model are for black-box inputs, or for different inputs whose representations share some data. For more details on these models, we refer to [22, 11].

1.2 In place and over-place quadratic algorithms

We recall here that classical, quadratic, algorithms for polynomial remaindering and triangular matrix operations can be performed in-place. For any ring \mathbb{D} we have for instance the following over-place algorithms for triangular matrix operations, given in Algorithm 1.

Algorithm 1 Over-place quadratic triangular matrix operations

(left: matrix-vector multiplication; right: triangular system solve)

Input: $U \in \mathbb{D}^{m \times m}$ upper triangular and $\vec{v} \in \mathbb{D}^m$.

Read-only: U .

Output: $\vec{v} \leftarrow U \cdot \vec{v}$.

```

1: for  $i = 1$  to  $m$  do
2:   for  $j = 1$  to  $i - 1$  do
3:      $v_j += U_{j,i} v_i$ 
4:   end for
5:    $v_i \leftarrow U_{i,i} v_i$ ;
6: end for

```

Output: $\vec{v} \leftarrow U^{-1} \cdot \vec{v}$.

```

1: for  $i = m$  down-to 1 do
2:   for  $j = m$  down-to  $i + 1$  do
3:      $v_i -= U_{i,j} v_j$ 
4:   end for
5:    $v_i \leftarrow U_{i,i}^{-1} v_i$ ;      {If  $U_{i,i} \in \mathbb{D}^*$ }
6: end for

```

For any vector $\vec{v} = [v_1, \dots, v_m]$, we associate the polynomial of degree $\leq m - 1$, with these coefficients: $V(X) = \sum_{i=0}^{m-1} v_{i+1} X^i$. We then give a quadratic in-place remaindering in Algorithm 2, as the classical long division.

Algorithm 2 In-place quadratic polynomial remainder

Input: $\vec{a} \in \mathbb{D}^N$, $\vec{b} \in \mathbb{D}^M$ with $b_M \in \mathbb{D}^*$.

Read-only: \vec{a}, \vec{b} .

Output: $R(X) = A(X) \bmod B(X)$.

```

1: if  $N < M$  then return  $R(X) = A(X) = \sum_{i=0}^{N-1} a_{i+1} X^i$ . end if
2: Let  $n = N - M + 1$  and  $\vec{r} = [a_{n+1}, \dots, a_N]$ ;
3: for  $i = n$  down-to 1 do
4:   Let  $\vec{r} = [a_i, r_1, \dots, r_{M-1}]$ ;
5:    $r_M \leftarrow r_M \cdot b_M^{-1}$ ;      {Local quotient is  $r_M \cdot b_M^{-1}$ }
6:    $\vec{r} -= r_M \cdot \vec{b}$ ;                {Except for last entry (of index  $M$ )}
7: end for
8: return  $R(X) = \sum_{i=0}^{M-2} r_{i+1} X^i$ .

```

1.3 Fast algorithms

More generally, if $\mathfrak{M}(M)$ is a submultiplicative complexity bound on an algorithm computing the multiplication of polynomials of degree M , we suppose in the following that there exists (not in-place) algorithms, see, e.g., [12, 15, 25], such that:

- $M \times M$ Toeplitz matrix-vector multiplication requires less than $\mathcal{O}(\mathfrak{M}(M))$ operations and $\mathcal{O}(M)$ extra space. More precisely, if $t \cdot M$ extra space is required, the matrix-vector multiplication requires less than $\lambda_t \mathfrak{M}(M)$ operations;
- $M \times M$ Triangular Toeplitz system solve requires also $\mathcal{O}(\mathfrak{M}(M))$ operations and $\mathcal{O}(M)$ extra space. More precisely, if $s \cdot M$ extra space is required, the system solving requires less than $\lambda_s \mathfrak{M}(M)$ operations.

Also, for sparse inputs, we suppose:

- $M \times M$ sparse Toeplitz (constant number of non-zero terms) matrix-vector multiplication requires less than $\mathcal{O}(M)$ operations and $\mathcal{O}(M)$ extra space;
- $M \times M$ Triangular sparse Toeplitz (constant number of non-zero terms) system solve requires less than $\mathcal{O}(M)$ operations and $\mathcal{O}(M)$ extra space.

2 Accumulated in-place and over-place Toeplitz operations

We thus now turn to in-place accumulated and over-place Toeplitz matrix-vector operations. Toeplitz matrix-vector multiplication can be reduced to circulant matrix-vector multiplication, via an embedding into a double-size circulant matrix. But this is not immediately in-place, since doubling the size requires a double space. We see in the following how we instead double the operations while keeping the same dimension. We start by the usual definitions, also extending circulant matrices to f-circulant matrices, following [23].

Definition 1. For $\vec{a} \in \mathbb{D}^m$, the circulant matrix represented by \vec{a} , $\mathcal{C}(\vec{a})$, is the $m \times m$ matrix (C_{ij}) , such that $C_{1j} = a_j$ and the $(i+1)$ -th row is the cyclic right shift by 1 of the i -th row.

Definition 2. For $f \in \mathbb{D}$ and $\vec{a} \in \mathbb{D}^m$, the f-circulant matrix represented by \vec{a} , $\mathcal{C}_f(\vec{a})$, is the $m \times m$ matrix (Γ_{ij}) , such that:

$$\text{for } C = \mathcal{C}(\vec{a}), \begin{cases} \Gamma_{ij} = C_{ij} & \text{if } i \leq j, \\ \Gamma_{ij} = f \cdot C_{ij} & \text{otherwise.} \end{cases}$$

Definition 3. For $\vec{a} \in \mathbb{D}^{2m-1}$, the (square) Toeplitz matrix represented by \vec{a} , $\mathcal{T}(\vec{a})$, is the $m \times m$ matrix (T_{ij}) , such that $T_{ij} = a_{m+j-i}$. Similarly, for

$\vec{a} \in \mathbb{D}^{m+n-1}$, we denote by $\mathcal{T}_{m,n}(\vec{a})$ the $m \times n$ rectangular Toeplitz matrix defined by its first column, $\vec{a}_{1..m}$ bottom to top, and its first row, $\vec{a}_{m..(m+n-1)}$, left to right.

We see immediately that with these notations, we have for instance $\mathcal{C}_1(\vec{a}) = \mathcal{C}(\vec{a})$, $\mathcal{C}_0(\vec{a}) = \frac{1}{2}(\mathcal{C}_1(\vec{a}) + \mathcal{C}_{-1}(\vec{a}))$, or also $\mathcal{C}_f(\vec{a}) = \mathcal{T}([f \cdot \vec{a}_{2..m}, \vec{a}])$, where $[\vec{u}, \vec{v}]$ denotes the vector obtained by concatenation of \vec{u} and \vec{v} .

It is well known that circulant matrices are diagonalized by a discrete Fourier transform, and hence can be manipulated via fast Fourier transforms as $\mathcal{C}_{\vec{a}} = F_m^{-1} \text{diag}(F_m \vec{a}) F_m$, for F_m a DFT-matrix, see, e.g., [13, § 4.7.7]. This gives a first way to compute circulant matrix-vector multiplication in-place (using and restoring afterwards both the matrix and the vector). Indeed, the (even truncated) Fourier transform and its inverse can be computed in-place [22, 14, 8] and in-place algorithm to compute $\vec{c} = \mathcal{C}_{\vec{a}} \cdot \vec{b}$ is:

1. $\vec{a} \leftarrow F_m \vec{a}$ and $\vec{b} \leftarrow F_m \vec{b}$;
2. $\vec{c} \leftarrow \text{diag}(\vec{a}) \cdot \vec{b}$;
3. $\vec{c} \leftarrow F_m^{-1} \vec{c}$; $\vec{b} \leftarrow F_m^{-1} \vec{b}$ and $\vec{a} \leftarrow F_m^{-1} \vec{a}$.

We use this idea and generalize it to f-circulant matrices, using fast methods for f-circulant matrices of [9, 23].

More precisely, we extend [23, Alg. 1] to (1) dimensions that are not a power of two (without any extra space for virtual zeroes), and (2) to any odd characteristic (i.e. we do not need an n -th root of unity), using techniques adapted from [10, Alg. 19].

To preserve a small memory footprint, it is crucial to handle carefully both cases where dimensions are not powers of two, and fields without roots of unity. On the one hand, using the closest larger power of two could double the required space. On the other hand, using an extended field would also multiply the size of the representation.

Our new algorithm is given as Algorithm 3. We use peeling to handle odd dimensions, and sum of squares to handle square roots. For now, we have two cases where this algorithm does not work:

- We need a division by 2, and thus cannot handle the even characteristic case;
- We need at least two squares, and thus cannot handle the field with 3 elements (our algorithm works over an extension, the field with 9 elements for instance, but this would, in theory, double the required space).

The cases we can tackle thus include any field of characteristic zero (complexes, reals, rationals, . . .), the characteristic 3 with more than 9 elements, any other odd characteristic.

Theorem 4. *Over any field with zero or odd characteristic (except the field of cardinality 3), Algorithm 3 is correct and requires less than $\mathcal{O}(m \log m)$ operations if f is a m -th power, and less than $\mathcal{O}(m \log^2 m)$ operations otherwise.*

Algorithm 3 IPfCmv: In-place accumulated f -circulant matrix-vector multiplication

Input: $f \in \mathbb{D}$, $\vec{a}, \vec{b}, \vec{c} \in \mathbb{D}^m$.

Read-only: f .

Output: $\vec{c} += \mathcal{C}_f(\vec{a}) \cdot \vec{b}$.

```

1: if  $m \leq \text{Threshold}$  then                                {Constant-time if Threshold  $\in \mathcal{O}(1)$ }
2:   return the quadratic in-place matrix-vector multiplication.
3: end if
4: if  $m$  is odd then
5:   Let  $\hat{a} = \vec{a}_{1..m-1}$ ,  $\hat{a} = \vec{a}_{2..m}$ ,  $\hat{a} = [a_m, \dots, a_1]$ ;
6:   Let  $\hat{b} = \vec{b}_{1..m-1}$ ,  $\hat{c} = \vec{c}_{1..m-1}$ ;
7:    $\hat{c} += \mathcal{C}_f(\hat{a}) \cdot \hat{b}$ ;                                {IPfCmv with even dimension}
8:    $\vec{c} += b_m \cdot \hat{a}$ ;                                    {Linear in-place accumulation}
9:    $c_m += f \cdot \hat{a} \cdot \hat{b}$ ;                            {Linear in-place dotproduct}
10: else if  $f$  is not a square, or  $f = 0$  then
11:   repeat
12:     Let  $\lambda \stackrel{\$}{\leftarrow} \mathbb{D}^*$ , s.t.  $(\lambda + 1) \in \mathbb{D}^*$ ; Let  $g = 1 + (\lambda + 1)(f - 1)$ ;
13:     until  $g$  is a square
14:      $\vec{c}^* = (1 + \lambda^{-1}) \cdot \vec{b}$ ;
15:      $\vec{c} += \mathcal{C}_1(a) \cdot \vec{b}$ ;                                {Direct in-place circulant call}
16:      $\vec{c}^* = \lambda \cdot \vec{c}^*$ ;
17:      $\vec{c} += \mathcal{C}_g(A) \cdot \vec{b}$ ;                                {IPfCmv call with a square}
18:      $\vec{c}^* = (\lambda + 1)^{-1} \cdot \vec{c}^*$ ;
19:   else
20:     Let  $a_1 = \vec{a}_{1..m/2}$ ,  $a_2 = \vec{a}_{m/2+1..m}$ ;
21:     Let  $b_1 = \vec{b}_{1..m/2}$ ,  $b_2 = \vec{b}_{m/2+1..m}$ ;                { $m$  is even}
22:     Let  $c_1 = \vec{c}_{1..m/2}$ ,  $c_2 = \vec{c}_{m/2+1..m}$ ;
23:      $\begin{cases} \bar{c}_1 \leftarrow c_1 \sqrt{f} + c_2 \\ \bar{c}_2 \leftarrow c_1 \sqrt{f} - c_2 \end{cases}$           {Linear in-place butterflies}
24:      $\begin{cases} \bar{b}_1 \leftarrow b_1 \sqrt{f} + b_2 \\ \bar{b}_2 \leftarrow b_1 \sqrt{f} - b_2 \end{cases}$           {Linear in-place butterflies}
25:      $\begin{cases} \bar{a}_1 \leftarrow a_1 + a_2 \sqrt{f} \\ \bar{a}_2 \leftarrow a_1 - a_2 \sqrt{f} \end{cases}$           {Linear in-place butterflies}
26:      $\bar{c}_1 += \mathcal{C}_{\sqrt{f}}(\bar{a}_1) \cdot \bar{b}_1$ ;                    {Recursive IPfCmv call}
27:      $\bar{c}_2 += \mathcal{C}_{-\sqrt{f}}(\bar{a}_2) \cdot \bar{b}_2$ ;                {Recursive IPfCmv call}
28:      $\begin{cases} a_1 \leftarrow (\bar{a}_1 + \bar{a}_2)/2 \\ a_2 \leftarrow (\bar{a}_1 - \bar{a}_2)/(2\sqrt{f}) \end{cases}$           {Linear in-place butterflies}
29:      $\begin{cases} b_1 \leftarrow (\bar{b}_1 + \bar{b}_2)/(2\sqrt{f}) \\ b_2 \leftarrow (\bar{b}_1 - \bar{b}_2)/2 \end{cases}$           {Linear in-place butterflies}
30:      $\begin{cases} c_1 \leftarrow (\bar{c}_1 + \bar{c}_2)/(2\sqrt{f}) \\ c_2 \leftarrow (\bar{c}_1 - \bar{c}_2)/2 \end{cases}$           {Linear in-place butterflies}
31:   end if
32: return  $\vec{c}$ .

```

Proof. When m is even and f is a square, the correctness is that of [23, Alg. 1]. When f is not a square, we decompose $f \neq 1$ as a sum of squares such that: $(\lambda + 1)f = \lambda^2 + b^2$. That is $b^2 = g = (\lambda + 1)(f - 1) + 1$. For any $b \neq 0$, $\lambda = (b^2 - 1)/(f - 1) - 1$ is a possibility if both λ and $\lambda + 1$ are non-zero. As soon as the field contains at least 3 distinct non-zero squares, there is always a solution. This includes fields of zero characteristic, and finite fields with at least 7 elements (since there are $(q - 1)/2$ nonzero squares in \mathbb{F}_q). In the field with 5 elements, both 1 and -1 are squares. It remains to remark that $\mathcal{C}_2 = -\mathcal{C}_1 + 2\mathcal{C}_{-1}$ and $\mathcal{C}_3 = 2\mathcal{C}_1 - \mathcal{C}_{-1}$. Thus only the case of \mathbb{F}_3 does not have a solution. Then,

$$\begin{aligned} & \left((\vec{c}(1 + \lambda^{-1}) + \mathcal{C}_1(A)\vec{b})\lambda + \mathcal{C}_g(A)\vec{b} \right) \frac{1}{1 + \lambda} \\ & = \vec{c} + \left(\frac{\lambda}{1 + \lambda} \mathcal{C}_1(A) + \frac{1}{1 + \lambda} \mathcal{C}_g(A) \right) \vec{b} \quad (1) \end{aligned}$$

Now, $\frac{\lambda}{1 + \lambda} + \frac{1}{1 + \lambda} = 1$ and $\frac{\lambda}{1 + \lambda} + \frac{1}{1 + \lambda}g = \frac{\lambda + 1 + (\lambda + 1)(f - 1)}{1 + \lambda} = f$, so that

$$\left(\frac{\lambda}{1 + \lambda} \mathcal{C}_1(A) + \frac{1}{1 + \lambda} \mathcal{C}_g(A) \right) = \mathcal{C}_f(A)$$

and the non square branch of the algorithm is also correct. Now for the odd dimension branch. We have that:

$$\mathcal{C}_f(\vec{a}) = \left[\begin{array}{ccc|c} a_1 & \dots & a_{m-1} & a_m \\ \cdot & \cdot & \cdot & \vdots \\ \hline f \cdot a_2 & \dots & f \cdot a_m & a_1 \end{array} \right].$$

Therefore the matrix-vector multiplication can be performed via $\mathcal{C}_f(\hat{a})$ for the top left corner, via a scalar multiplication for the right-most column and via a dot-product for the bottom left row. For being in-place, each of the 3 vectors \vec{a} , \vec{b} , \vec{c} need initial and final linear combinations of halves. Those can be performed sequentially with, say, two temporaries: this is $\mathcal{O}(n)$ classical butterfly swaps of the generic form $t = a_i r$; $u = b_i s$; $\bar{a}_i = t + u$; $\bar{b}_i = t - u$; $\bar{a}_i * = v$; $\bar{b}_i * = w$. Now for the complexity bound. When called with an m -th power (for instance $f = 1$) the algorithm will always take the branch where f is a square, and therefore requires $T(m) = 6m + 2T(m/2) + 6m$ operations, that is $\mathcal{O}(m \log m)$. Now when called with anything else, in the worst case it always takes the non-square branch. Then the cost is $T(m) = m + \mathcal{O}(m \log m) + m + (2T(m/2) + 12m) + m$, so that it becomes $\mathcal{O}(m \log^2 m)$. These bounds also bound the odd cases as the smallest power of two larger than x is $\mathcal{O}(x)$. \square

Theorem 4 does not work in particular in characteristic 2, as it divides by 2. We show how to deal with most of the characteristic 2 cases in Appendix B.

Remark 5. We here present f -circulant matrices where the coefficient acts on the lower left part of the matrix (excluding the diagonal). In the same manner, one can design an in-place fast algorithm for the other type of f -circulant

matrices, where the coefficient would act on the upper right part of the matrix (excluding the diagonal).

This generalization of fast algorithms for f-circulant matrices allows us to build algorithms for accumulation with triangular Toeplitz matrices first, as sums of f-circulant in Algorithms 4 and 5, and then for any Toeplitz matrix, again as sums of triangular Toeplitz matrices in Algorithm 6.

Algorithm 4 In-place accumulated Upp. Triang. Toeplitz m-v. mult.

Input: $\vec{a}, \vec{b}, \vec{c} \in \mathbb{D}^m$.

Output: $\vec{c} += \mathcal{T}([\vec{0}, \vec{a}]) \cdot \vec{b}$.

1: **return** $\vec{c} += \mathcal{C}_0(\vec{a}) \cdot \vec{b}$. {Algorithm 3}

Algorithm 5 In-place accumulated Low. Triang. Toeplitz m-v. mult.

Input: $\vec{a} \in \mathbb{D}^{m-1}, \vec{b}, \vec{c} \in \mathbb{D}^m$.

Output: $\vec{c} += \mathcal{T}([\vec{a}, \vec{0}]) \cdot \vec{b}$.

1: $\vec{c} += \mathcal{C}_2([0, \vec{a}]) \cdot \vec{b}$; {Algorithm 3}

2: **return** $\vec{c} -= \mathcal{C}_1([0, \vec{a}]) \cdot \vec{b}$. {Algorithm 3}

Remark 6. Algorithm 4 deals with upper triangular Toeplitz matrices with invertible diagonal element and Algorithm 5 with lower triangular Toeplitz matrices with zero diagonal element. From Theorem 5, one could likewise derive the transposed versions and this would provide instead an algorithm for lower triangular Toeplitz matrices with a single call to Algorithm 3.

Algorithm 6 In-place accumulated square Toeplitz m-v. mult.

Input: $\vec{a}_1 \in \mathbb{D}^{m-1}, \vec{a}_2, \vec{b}, \vec{c} \in \mathbb{D}^m$,

Output: $\vec{c} += \mathcal{T}([\vec{a}_1, \vec{a}_2]) \cdot \vec{b}$.

1: $\vec{a}_2 -= [0, \vec{a}_1]$;

2: $\vec{c} += \mathcal{C}_0(\vec{a}_2) \cdot \vec{b}$; {Algorithm 3}

3: $\vec{a}_2 += [0, \vec{a}_1]$;

4: **return** $\vec{c} += \mathcal{C}_1([0, \vec{a}_1]) \cdot \vec{b}$. {Algorithm 3}

Lemma 7. Algorithms 4 to 6 are correct.

Proof. Directly looking at the matrices values, we see that:

- $\mathcal{T}([\vec{0}, \vec{a}]) = \mathcal{C}_0(\vec{a})$.
- $\mathcal{T}([\vec{a}, \vec{0}]) = \mathcal{C}_2([\vec{a}, 0]) - \mathcal{C}_1([\vec{a}, 0])$.
- $\mathcal{T}([\vec{a}_1, \vec{a}_2]) = \mathcal{C}_0(\vec{a}_2 - [0, \vec{a}_1]) + \mathcal{C}_1([0, \vec{a}_1])$ [21, Eq. (2.6.2)].

□

From this, we give in Algorithm 7 an in-place rectangular Toeplitz matrix-vector multiplication.

Algorithm 7 In-place accumulated rectangular Toeplitz m-v. mult.

Input: $\vec{a} \in \mathbb{D}^{m+n-1}$, $\vec{b} \in \mathbb{D}^n$, $\vec{c} \in \mathbb{D}^m$,

Output: $\vec{c} += \mathcal{T}_{m,n}(\vec{a}) \cdot \vec{b}$.

```

1: if  $m = n$  then
2:   return  $\vec{c} += \mathcal{T}(\vec{a}) \cdot \vec{b}$ . {Algorithm 6}
3: end if
4: if  $m > n$  then
5:   Let  $c_1 = \vec{c}_{1..n}$  and  $c_2 = \vec{c}_{(n+1)..m}$ ;
6:    $c_1 += \mathcal{T}(\vec{a}_{(m-n+1)..(m+n-1)}) \cdot \vec{b}$ ; {Algorithm 6}
7:    $c_2 += \mathcal{T}_{m-n,n}(\vec{a}_{1..(m-1)}) \cdot \vec{b}$ ; {Recursive call}
8: else
9:   Let  $b_1 = \vec{b}_{1..m}$  and  $b_2 = \vec{b}_{(m+1)..n}$ ;
10:   $c += \mathcal{T}(\vec{a}_{1..(2m-1)}) \cdot b_1$ ; {Algorithm 6}
11:   $c += \mathcal{T}_{m,n-m}(\vec{a}_{(m+1)..(m+n-1)}) \cdot b_2$ ; {Recursive call}
12: end if
13: return  $\vec{c}$ .
```

Proposition 8. *Algorithm 7 is correct and requires less than*

$$\mathcal{O}(\max\{m, n\} \log^2(\min\{m, n\}))$$

operations.

Proof. If $m > n$ then there are first $\lfloor m/n \rfloor$ square $n \times n$ calls. This is bounded by $\mathcal{O}((m/n)n \log^2(n))$ operations. The recursive call is then negligible. This is similar when $m < n$. □

These in-place accumulated Toeplitz matrix-vector multiplications, in turns, allows us to obtain an over-place triangular Toeplitz multiplication or system solve, given in Algorithms 8 and 9.

Proposition 9. *Algorithm 8 is correct and requires less than $\mathcal{O}(m \log^3 m)$ operations.*

Proof. For the correctness, let $T = \mathcal{T}([\vec{a}, \vec{0}])$ and consider it as blocks $T_1 = \mathcal{T}([a_{(m-k+1)..m}, \vec{0}])$, $T_2 = \mathcal{T}([\vec{a}_{(k+1)..m}, \vec{0}])$ and $G = \mathcal{T}_{m-k,k}([a_1, \dots, a_{m-1}])$. Then $T = \begin{pmatrix} T_1 & 0 \\ G & T_2 \end{pmatrix}$. Thus $T\vec{b} = \begin{pmatrix} T_1 b_1 \\ G b_1 + T_2 b_2 \end{pmatrix}$. Let $\bar{b}_1 = T_1 b_1$, $\hat{b}_2 = T_2 b_2$ and $\bar{b}_2 = G b_1 + T_2 b_2$. Then $\bar{b}_2 = \hat{b}_2 + G b_1$ and the algorithm is correct.

Now for the complexity bound, the cost function is

$$T(m) \leq 2T(m/2) + \mathcal{O}(m \log^2 m) = \mathcal{O}(m \log^3 m).$$

□

Algorithm 8 Over-place triang. Toeplitz m-v. mult.

Input: $\vec{a}, \vec{b} \in \mathbb{D}^m$, s.t. $a_1 \in \mathbb{D}^*$.

Output: $\vec{b} \leftarrow \mathcal{T}(\vec{a}, \vec{0}) \cdot \vec{b}$.

- 1: **if** $m \leq \text{Threshold}$ **then** {Constant-time if Threshold $\in \mathcal{O}(1)$ }
 - 2: **return** the quadratic in-place triangular m-v. mult.
 - 3: **end if**
 - 4: Let $k = \lceil m/2 \rceil$, $b_1 = \vec{b}_{1..k}$ and $b_2 = \vec{b}_{(k+1)..m}$;
 - 5: $b_2 \leftarrow \mathcal{T}(\vec{a}_{(k+1)..m}, \vec{0}) \cdot b_2$; {Recursive call}
 - 6: $b_2 \text{ += } \mathcal{T}_{m-k,k}([a_1, \dots, a_{m-1}]) \cdot b_1$; {Algorithm 7}
 - 7: $b_1 \leftarrow \mathcal{T}(a_{(m-k+1)..m}, \vec{0}) \cdot b_1$; {Recursive call}
 - 8: **return** \vec{b} .
-

Algorithm 9 Over-place triang. Toeplitz system solve

Input: $\vec{a}, \vec{b} \in \mathbb{D}^m$, s.t. $a_1 \in \mathbb{D}^*$.

Output: $\vec{b} \leftarrow \mathcal{T}(\vec{0}, \vec{a})^{-1} \cdot \vec{b}$.

- 1: **if** $m \leq \text{Threshold}$ **then** {Constant-time if Threshold $\in \mathcal{O}(1)$ }
 - 2: **return** the quadratic in-place triangular system solve.
 - 3: **end if**
 - 4: Let $k = \lceil m/2 \rceil$, $b_1 = \vec{b}_{1..k}$ and $b_2 = \vec{b}_{(k+1)..m}$;
 - 5: $b_2 \leftarrow \mathcal{T}(\vec{0}, \vec{a}_{(k+1)..m})^{-1} \cdot b_2$; {Recursive call}
 - 6: $b_1 \text{ -= } \mathcal{T}_{k,m-k}(\vec{a}_{2..m}) \cdot b_2$; {Algorithm 7}
 - 7: $b_1 \leftarrow \mathcal{T}(\vec{0}, \vec{a}_{1..k})^{-1} \cdot b_1$; {Recursive call}
 - 8: **return** \vec{b} .
-

Proposition 10. *Algorithm 9 is correct and requires less than $\mathcal{O}(m \log^3 m)$ operations.*

Proof. For the correctness, let $T = \mathcal{T}([\vec{0}, \vec{a}])$ and consider it as blocks $T_1 = \mathcal{T}([\vec{0}, \vec{a}_{1..k}])$, $T_2 = \mathcal{T}([\vec{0}, \vec{a}_{(k+1)..m}])$ and $G = \mathcal{T}_{k,m-k}(\vec{a}_{2..m})$. Then $T = \begin{pmatrix} T_1 & G \\ 0 & T_2 \end{pmatrix}$. Now define H , s.t. $T^{-1} = \begin{pmatrix} T_1^{-1} & H \\ 0 & T_2^{-1} \end{pmatrix}$. Then H satisfies $T_1^{-1}G + HT_2 = 0$. Also, we have $T^{-1}\vec{b} = \begin{pmatrix} T_1^{-1}b_1 + Hb_2 \\ T_2^{-1}b_2 \end{pmatrix}$. Let $\bar{b}_2 = T_2^{-1}b_2$ and $\bar{b}_1 = T_1^{-1}b_1 + Hb_2$. Then $\bar{b}_1 = T_1^{-1}b_1 + HT_2\bar{b}_2 = T_1^{-1}b_1 - T_1^{-1}G\bar{b}_2 = T_1^{-1}(b_1 - G\bar{b}_2)$.
Now for the complexity bound, the cost function is

$$T(m) \leq 2T(m/2) + \mathcal{O}(m \log^2 m) = \mathcal{O}(m \log^3 m).$$

□

We now have new over-place Toeplitz methods. Next, we start by reducing the extra space for polynomial remaindering. Eventually, we will combine these two techniques to obtain in-place polynomial remaindering.

3 Over-writing the quotient

With two polynomials A and B of respective degrees N and M , the computation of the euclidean division remainder R of degree strictly less than M such that $A = BQ + R$ with quotient Q , can be rewritten as $R \equiv A - BQ \pmod{X^M}$. This is therefore enough to compute the quotient only up to the degree $M - 1$: let $A_M \equiv A \pmod{X^M}$ and $Q_M \equiv Q \pmod{X^M}$, then $R \equiv A_M - BQ_M \pmod{X^M}$.

This observation is the ingredient that allows to compute the remainder using an extra space only of the order of the degree of the divisor B . One can also see this as the long division algorithm applied to blocks of dimension M .

Let us write the euclidean equation $A = BQ + R$ in a Toeplitz matrix form. We view the polynomials A , Q and R as vectors $[a_0, \dots, a_N]$, $[q_0, \dots, q_{N-M}]$, $[r_0, \dots, r_{M-1}, 0, \dots, 0]$ and B as a Toeplitz matrix $B = \mathcal{T}([0, \dots, 0, b_M, \dots, b_0, 0, \dots, 0])$ as follows:

$$\begin{bmatrix} a_0 \\ \vdots \\ \vdots \\ \vdots \\ a_N \end{bmatrix} = \begin{bmatrix} b_0 & & & 0 \\ \vdots & \ddots & & \\ b_M & \ddots & b_0 & \\ & \ddots & \vdots & \\ 0 & & b_M & \end{bmatrix} \cdot \begin{bmatrix} q_0 \\ \vdots \\ \vdots \\ q_{N-M} \end{bmatrix} + \begin{bmatrix} r_0 \\ \vdots \\ r_{M-1} \\ 0 \end{bmatrix}. \quad (2)$$

Focusing on the last $N - M + 1$ rows of Eq. (2) we obtain directly the upper triangular $(N - M + 1) \times (N - M + 1)$ Toeplitz system of equations whose solution is only the quotient:

$$\begin{bmatrix} a_M \\ \vdots \\ \vdots \\ a_N \end{bmatrix} = \begin{bmatrix} b_M & \dots & b_0 & & 0 \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & \\ & & & \ddots & \\ 0 & & & & b_M \end{bmatrix} \cdot \begin{bmatrix} q_0 \\ \vdots \\ \vdots \\ q_{N-M} \end{bmatrix}. \quad (3)$$

We now let $n = N - M + 1$ and suppose that B is really of degree M , that is, its leading coefficients b_M is invertible in the coefficient domain. For the sake of simplicity we also assume that n is a multiple of M (otherwise just complete the polynomial A with virtual leading zero coefficients up to the next multiple of M) and let $\mu = \frac{n}{M}$. We then denote the $M \times M$ blocks of the Toeplitz matrix in Eq. (3) by: $T = \mathcal{T}([\vec{0}_{M-1}, b_M, \dots, b_1])$ and $G = \mathcal{T}([b_{M-1}, \dots, b_0, \vec{0}_{M-1}])$, that is:

$$T = \begin{bmatrix} b_M & \dots & b_1 \\ & \ddots & \vdots \\ & & b_M \end{bmatrix} \text{ and } G = \begin{bmatrix} b_0 & & \\ \vdots & \ddots & \\ b_{M-1} & \dots & b_0 \end{bmatrix}. \quad (4)$$

This in turns gives a way to access only the first coefficients of Q in an upper triangular Toeplitz system, with a 2-block band structure:

$$\begin{bmatrix} q_0 \\ \vdots \\ \vdots \\ q_{M-1} \end{bmatrix} = \begin{bmatrix} I_M & \vdots & 0 \end{bmatrix} \begin{bmatrix} T & G & & 0 \\ & \ddots & \ddots & \\ & & \ddots & G \\ 0 & & & T \end{bmatrix}^{-1} \begin{bmatrix} a_M \\ \vdots \\ \vdots \\ a_N \end{bmatrix} \quad (5)$$

where $\begin{bmatrix} I_M & \vdots & 0 \end{bmatrix}$ is the concatenation of the $M \times M$ identity matrix and the $(N - M) \times (N - M)$ zero matrix.

Finally, recovering the remainder from the first M rows of Eq. (2) equations is just like multiplying the quotient by G and thus $R = A - BQ \pmod{X^M}$ can be written as:

$$R = \begin{bmatrix} a_0 \\ \vdots \\ \vdots \\ a_{M-1} \end{bmatrix} - \begin{bmatrix} G & \vdots & 0 \end{bmatrix} \begin{bmatrix} T & G & & 0 \\ & \ddots & \ddots & \\ & & \ddots & G \\ 0 & & & T \end{bmatrix}^{-1} \begin{bmatrix} a_M \\ \vdots \\ \vdots \\ a_N \end{bmatrix} \quad (6)$$

In fact, as we need only the first M coefficient of the Toeplitz system we just need the first block-row of the (upper triangular) inverse of the upper triangular 2-band Toeplitz matrix.

Now this first block-row of the inverse of an upper triangular Toeplitz matrix U is given by a direct formula obtained from either of the equations $U \cdot U^{-1} = I$

or $U^{-1} \cdot U = I$ (see, e.g., [7, Eq.(1)] for the scalar case). If we denote by H_i the blocks of that row, we have that:

$$\begin{cases} H_1 = T^{-1} \\ H_{i-1}G + H_iT = 0, & i = 2..\mu \\ TH_i + GH_{i-1} = 0, & i = 2..\mu \end{cases} \quad (7)$$

Solving Eq. (7), we just get that $H_i = T^{-1}(-GT^{-1})^i = (-GT^{-1})^iT^{-1}$.

We have shown:

Lemma 11.

$$\begin{bmatrix} T & G & & 0 \\ & \ddots & \ddots & \\ & & \ddots & G \\ 0 & & & T \end{bmatrix}^{-1} = T^{-1} \cdot \begin{bmatrix} I & -GT^{-1} & \dots & (-GT^{-1})^{\mu-1} \\ & \ddots & \ddots & \vdots \\ & & \ddots & -GT^{-1} \\ 0 & & & I \end{bmatrix}$$

Now denote by $[\vec{a}_0, \vec{a}_1, \dots, \vec{a}_\mu]$ the decomposition into blocks of dimension M of $[a_0, \dots, a_N]$. Combining Eq. (6) and Theorem 11, we obtain now that:

$$R = \begin{bmatrix} a_0 \\ \vdots \\ \vdots \\ a_{M-1} \end{bmatrix} - GT^{-1} \left(\sum_{i=1}^{\mu} (-GT^{-1})^{i-1} \vec{a}_i \right) = \sum_{i=0}^{\mu} (-GT^{-1})^i \vec{a}_i \quad (8)$$

From Eq. (8) we thus can immediately deduce the following Algorithm 10 that uses only $\mathcal{O}(M)$ extra memory space in a Horner-like fashion of the polynomial in $(-GT^{-1})$ of Eq. (8). Note that this algorithm does not modify its input along its course: both $A(X)$ and $B(X)$ are for now read-only.

Theorem 12. *Algorithm 10 is correct and requires $\mathcal{O}(\frac{N}{M}\mathfrak{M}(M))$ arithmetic operations and $\mathcal{O}(M)$ extra memory space. If the polynomial B is sparse with a constant number of non-zero coefficients, the arithmetic complexity is reduced to $\mathcal{O}(N)$.*

Proof. Correctness stems directly from Theorem 11 and Eq. (8). For the complexity bounds, we use the results of Section 1.3. For each block, the triangular Toeplitz system solve and the Toeplitz m-v. mult. require respectively $\lambda_s\mathfrak{M}(m)$ and $\lambda_t\mathfrak{M}(m)$ operations and, sequentially, $\max\{s; t\}M$ extra space. Apart from this space, we only need one extra vector, \vec{t} , to store intermediate results. Overall we thus perform $\mu((\lambda_s + \lambda_t)\mathfrak{M}(M) + M)$ operations. With $\mu = n/M$ and $n = N - M + 1$, this is $\lceil \frac{N-M+1}{M} \rceil ((\lambda_s + \lambda_t)\mathfrak{M}(M) + M) = \mathcal{O}(\frac{N}{M}\mathfrak{M}(M))$ operations, using $(1 + \max\{s; t\})M$ extra space.

Now if B is sparse with a constant number of non-zero elements, each triangular Toeplitz system solve and Toeplitz matrix-vector multiplication can be performed with only $\mathcal{O}(M)$ operations with the same extra memory space. Thus the overall arithmetic bound becomes $\mathcal{O}(\mu M) = \mathcal{O}(N)$. \square

Algorithm 10 Overwritten-Quotient euclidean remainder

Input: $A(X), B(X)$ in $\mathbb{D}[X]$ of respective degrees N and M .

Read-only: $A(X), B(X)$.

Output: $R(X) \equiv A(X) \pmod{B(X)}$ of degree at most $M - 1$.

```
1: if  $M > N$  then return  $A$ .
2: Let  $n = N - M + 1, \mu = \lceil \frac{n}{M} \rceil$ ;
3: Let  $[\vec{a}_0, \dots, \vec{a}_\mu] = [a_0, \dots, a_N, \vec{0}]$ ;                                {Blocks of dimension  $M$ }
4: Let  $T = \mathcal{T}([\vec{0}_{M-1}, b_M, \dots, b_1]), G = \mathcal{T}([b_{M-1}, \dots, b_0, \vec{0}_{M-1}])$ ;
5:  $\vec{r} = \vec{a}_\mu$ ;                                                                { $\vec{r}$  in-place of the result}
6: for  $i = \mu - 1$  down-to 0 do
7:    $\vec{t} = T^{-1} \cdot \vec{r}$ ;                                                        {Triang. Toeplitz solve}
8:    $\vec{r} = (-G) \cdot \vec{t}$ ;                                                       {Triang. Toeplitz m-v. mult.}
9:    $\vec{r} += \vec{a}_i$ ;
10: end for
11: return  $R = \sum_{i=0}^{M-1} r_i X^i$ .
```

Remark 13. *This Algorithm 10 is in fact just the long division polynomial algorithm applied to sub-blocks of the polynomial of size M :*

- $\vec{t} = T^{-1} \cdot \vec{r}$, Line 7, corresponds to computing the quotient of the current leading coefficients, of the dividend, by B ;
- $\vec{r} = (-G) \cdot \vec{t}$, Line 8, corresponds to recovering the lower part of the multiplication of that current quotient by B ;
- $\vec{r} += \vec{a}_i$, Line 9, updates the next M coefficients of the current dividend (the leading ones being zero by construction of the current quotient).

This in-place long division by block is also sketched for instance in the proof of [12, Lemma 2.1]. The latter Lemma gives about $3\lambda\mathfrak{M}(N)$ operations and $(2 + s)M$ extra space. In fact a refined analysis should also give the same (better) complexity as Theorem 12, that is less than $2\lambda\frac{N}{M}\mathfrak{M}(M)$ operations and $(1 + s)M$ extra space.

4 Time-memory trade-off with more bands

Here we present a variant that enables to balance speed with storage. The algorithm still uses blocks, but requires only a small fraction of extra space. If the output space is M , then one obtains an algorithm using the space of one of its input, M space for the output plus only a fraction $m = M/\mu$ of extra space. The drawback is that our complexity bounds become $\mathcal{O}(\frac{N}{m}\frac{M}{m}\mathfrak{M}(m))$. Therefore this opens a whole range of algorithms in practice, from $m = \mathcal{O}(1)$, that is a constant extra space but quadratic complexity $\mathcal{O}(NM)$; to $\mu = \mathcal{O}(1)$, with $M(1 + o(1))$ total space required (including the output of size M , that is, the extra space can be a very small fraction of the output, but not quite a constant), and a complexity that remains $\mathcal{O}(\frac{N}{M}\mathfrak{M}(M))$.

Following [7, 24, 18, 17, 1, 19], we start by generalizing Theorem 11 to more Toeplitz bands:

Lemma 14. For $\alpha - 1$ matrices G_2, \dots, G_α of dimension $m \times m$ and for T_1 an $m \times m$ upper triangular matrix, consider the following inverse of the $(m\eta) \times (m\eta)$ upper triangular Toeplitz matrix with α non-zero blocks bands:

$$\begin{bmatrix} H_1 & \dots & H_\eta \\ & \ddots & \ddots \\ & & H_1 \end{bmatrix} = \begin{bmatrix} T_1 & G_2 & \dots & G_\alpha & & 0 \\ & \ddots & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & \ddots & G_\alpha \\ & & & \ddots & \ddots & \vdots \\ & & & & \ddots & G_2 \\ & & & & & T_1 \end{bmatrix}^{-1}.$$

Then

$$\begin{cases} H_1 = T_1^{-1} \\ \sum_{j=2}^{\min\{\alpha; i\}} H_{i-j+1} G_j + H_i T_1 = 0, & i = 2.. \eta, \\ T_1 H_i + \sum_{j=2}^{\min\{\alpha; i\}} H_{i-j+1} G_j = 0, & i = 2.. \eta. \end{cases} \quad (9)$$

Now we consider m consecutive coefficients of R in Eq. (6) for $m < M$. Let $\eta = \lceil \frac{N-M+1}{m} \rceil$ and $\alpha = \lceil \frac{M}{m} \rceil + 1$. If the consecutive coefficients start at index i_0 , let $k = \lceil \frac{i_0-1}{m} \rceil$. Then now let $\vec{\rho}_k = [a_{i_0}, \dots, a_{i_0+m-1}]$ and $[\vec{a}_1, \dots, \vec{a}_\eta] = [a_M, \dots, a_N, 0, \dots, 0]$ again completed with zeroes if $N - M + 1$ is not a multiple of m .

Let the first m rows of $[T \ G]$ be

$$[T_1 \ G_2 \ \dots \ G_\alpha \ 0] = [I_m \ \vdots \ 0] [T \ G] \quad (10)$$

Finally, consider the m consecutive rows of $[G \ \vdots \ 0]$ starting at row i_0 as $[F_0 \ \dots \ F_k \ \vdots \ 0]$

Then combining Eq. (6) and Theorem 14, we have the following:

Lemma 15. For A and B polynomials of respective degrees N and M we let $R \equiv A \pmod{B}$. Let $\vec{\rho}_k = [a_{i_0}, \dots, a_{i_0+m-1}]$ and $[\vec{a}_1, \dots, \vec{a}_\eta] = [a_M, \dots, a_N, 0, \dots, 0]$

and consider the Toeplitz blocks $T = \begin{bmatrix} b_M & \dots & b_1 \\ & \ddots & \vdots \\ & & b_M \end{bmatrix}$, $G = \begin{bmatrix} b_0 & & \\ \vdots & \ddots & \\ b_{M-1} & \dots & b_0 \end{bmatrix}$

together with the subsets of m rows $[T_1 \ G_2 \ \dots \ G_\alpha \ 0] = [I_m \ \vdots \ 0] [T \ G]$ and $[F_0 \ \dots \ F_k \ \vdots \ 0] = [0_{i_0} \ \vdots \ I_m \ \vdots \ 0_{M-i_0-m}] \cdot [G \ \vdots \ 0]$. Finally let $[H_1, \dots, H_\eta]$ be the components of the Toeplitz inverse as in Theorem 15, then:

$$\begin{bmatrix} r_{i_0} \\ \vdots \\ r_{i_0+m-1} \end{bmatrix} = \vec{\rho}_k - \sum_{j=0}^k F_j \left(\sum_{i=1}^{\eta-j} H_i \vec{a}_{i+j} \right) \quad (11)$$

This Theorem 15 together with the characterization of the inverse components of Eq. (9), now give us a way to compute this subset of the remainder via a generalized Horner-like process.

To save space and time the idea is to avoid the computation of the inverse altogether. For this, we recursively replace any computation of the form $H_i \vec{a}_\ell$, using Eq. (9) as follows: $H_i \vec{a}_\ell = H_i T_1 T_1^{-1} \vec{a}_\ell = (-\sum_{j=2}^{\min\{\alpha; i\}} H_{i-j+1} G_j) T_1^{-1} \vec{a}_\ell = \sum_{j=2}^{\min\{\alpha; i\}} H_{i-j+1} (-G_j T_1^{-1} \vec{a}_\ell)$.

The core computations are now again of the form $G_j T_1^{-1} \vec{a}_\ell$ as in Algorithm 10 but α of them have to be performed at each step. In order to save the space to store them (which would be $m\alpha$, that is not smaller than $mM/m = M$), the idea is to modify instead the appropriate \vec{a}_{i-j+1} parts of the input. The trick is that it is possible to do this up to $i = 1$, then store the overall sum into the final result space, and finally revert the whole computation (thus doubling the cost) to put back the input in its initial state. This is shown in Algorithm 11.

Theorem 16. *Algorithm 11 is correct and if Triangular Toeplitz solve requires $\lambda_s \mathfrak{M}(m)$ operations with $s \cdot m$ extra space and Toeplitz matrix-vector multiplication requires $\lambda_t \mathfrak{M}(m)$ operations with $t \cdot m$ extra space, then with $(2 + \max\{s; t\})m$ extra space, its arithmetic complexity is bounded by $\mathcal{O}(\frac{N}{m} \frac{M}{m} \mathfrak{M}(m))$.*

Proof. Correctness directly comes from Theorems 14 and 15. Then, the complexity bound is twice $\eta(\lambda_s \mathfrak{M}(m) + (\alpha - 1)\lambda_t \mathfrak{M}(m) + m)$ plus the middle computations $\sum_{k=0}^{\alpha-1} (k+1)(\lambda_s \mathfrak{M}(m) + \lambda_t \mathfrak{M}(m) + m)$. Its arithmetic complexity is thus bounded by:

$$\left(2\eta(\lambda_s + (\alpha - 1)\lambda_t) + \alpha \frac{\alpha + 1}{2} (\lambda_s + \lambda_t) \right) \mathfrak{M}(m) + \left(2\eta + \alpha \frac{\alpha + 1}{2} \right) m.$$

With $\eta = \lceil \frac{N-M+1}{m} \rceil$ and $\alpha = \lceil \frac{M}{m} \rceil + 1$ we obtain the claimed bound.

Now, if the operations are performed sequentially, then apart from the used input, \vec{a}_i , two size m memory slots are required, \vec{r} and \vec{t} , and either the extra space of the Toeplitz solving or that of the multiplication. This is $(2 + \max\{s; t\})m$ extra space. \square

Remark 17. *Again, we see that $\vec{r} = T_1^{-1} \cdot \vec{a}_i$ is computing the current quotient $\vec{r} = \text{Quo}(\vec{a}_i(X)X^{m-1}, B_1(X))$ and that the loop $(\vec{t} = G_j \cdot \vec{r})_{j=2.. \alpha}$ is the multiplication by the whole divisor, $B(X)\vec{r}(X)$.*

5 In-place modular remainder

We now derive algorithms that uses only $\mathcal{O}(1)$ extra memory space in the in-place model of Section 1.1: modifying the inputs is possible if and only if all inputs are restored to their initial state after the completion of the algorithm.

This allows us to store some intermediate results, over-writing the input, provided that we can afterwards recompute the initial inputs in their entirety. The idea is to combine Sections 2 to 4.

Algorithm 11 BRem: block remainder using & restoring its LHS

Input: $m, \vec{a} \in \mathbb{D}^N, \vec{b} \in \mathbb{D}^M, \eta = \lceil \frac{N-M+1}{m} \rceil, \alpha = \lceil \frac{M}{m} \rceil + 1;$

Read-only: $\vec{b}.$

Output: $R = [R_0, \dots, R_{\alpha-1}]$, with $R_k = \vec{\rho}_k - \sum_{j=0}^k F_j \left(\sum_{i=1}^{\eta-j} H_i \vec{a}_{i+j} \right)$ and the H_i as defined in Theorem 14.

```

1: Let  $\eta = \lceil \frac{N-M+1}{m} \rceil, \alpha = \lceil \frac{M}{m} \rceil + 1;$ 
2: Let  $[\vec{\rho}_0, \dots, \vec{\rho}_{\alpha-1}] = [a_0, \dots, a_{M-1}, 0, \dots, 0]$  and  $[\vec{a}_1, \dots, \vec{a}_\eta] = [a_M, \dots, a_N, 0, \dots, 0]$ 
3: Let  $T = \mathcal{T}([\vec{0}_{M-1}, b_M, \dots, b_1]), G = \mathcal{T}([b_{M-1}, \dots, b_0, \vec{0}_{M-1}])$ 
4: Let  $[T_1 \ G_2 \ \dots \ G_\alpha \ 0] = [I_m \ 0]$   $[T \ G]$ 
5: for  $i = \eta$  down-to 1 do                                     {Propagate  $\eta$  times Eq. (9) to  $\vec{a}$ }
6:    $\vec{r} = T_1^{-1} \cdot \vec{a}_i;$                                        {triangular Toeplitz solve}
7:   for  $j = 2$  to  $\min\{\alpha; i\}$  do
8:      $\vec{t} = G_j \cdot \vec{r};$                                            {Toeplitz matrix-vector multiplication}
9:      $\vec{a}_{i-j+1} -= \vec{t};$ 
10:  end for
11: end for
12: for  $k = 0$  to  $\alpha - 1$  do
13:   Let  $i_0 = k \cdot m;$ 
14:   Let  $[F_0 \ \dots \ F_k \ 0] = [0_{i_0} \ I_m \ 0_{M-i_0-m}] \cdot [G \ 0];$ 
15:    $R_k = \vec{\rho}_k;$ 
16:   for  $j = 0$  to  $k$  do                                           {Now this is just  $\vec{\rho}_k - \sum_{j=0}^k F_j T_1^{-1} \vec{a}_j$ }
17:      $\vec{r} = T_1^{-1} \cdot \vec{a}_j;$                                        {triangular Toeplitz solve}
18:      $\vec{t} = F_j \cdot \vec{r};$                                            {Toeplitz matrix-vector multiplication}
19:      $R_k -= \vec{t};$ 
20:  end for
21: end for
22: for  $i = 1$  to  $\eta$  do                                             {Unroll  $\eta$  times Eq. (9) to  $\vec{a}$ }
23:    $\vec{r} = T_1^{-1} \cdot \vec{a}_i;$                                        {triangular Toeplitz solve}
24:   for  $j = 2$  to  $\min\{\alpha; i\}$  do
25:      $\vec{t} = G_j \cdot \vec{r};$                                            {Toeplitz matrix-vector multiplication}
26:      $\vec{a}_{i-j+1} += \vec{t};$ 
27:  end for
28: end for
29: return  $R = [R_0, \dots, R_{\alpha-1}].$ 

```

First we present in Algorithm 12 a fully in-place remainder where only $B(X)$ is modified but restored: this variant replaces only Lines 7 to 9 of Algorithm 10.

Algorithm 12 In-place euclidean remainder (fully in-place variant of Algorithm 10)

Input: $A(X), B(X)$ in $\mathbb{D}[X]$ of respective degrees N and M .

Read-only: $A(X)$.

Output: $R(X) \equiv A(X) \pmod{B(X)}$ of degree at most $M - 1$.

```

1: if  $M > N$  then return  $A$ .
2: Let  $n = N - M + 1, \mu = \lceil \frac{n}{M} \rceil$ ;
3: Let  $[\vec{a}_0, \dots, \vec{a}_\mu] = [a_0, \dots, a_N, \vec{0}]$ ;                                {Blocks of dimension  $M$ }
4: Let  $T = \mathcal{T}([\vec{0}_{M-1}, b_M, \dots, b_1]), G = \mathcal{T}([b_{M-1}, \dots, b_0, \vec{0}_{M-1}])$ ;
5:  $\vec{r} = \vec{a}_\mu$ ;                                                                { $\vec{r}$  in-place of the result}
6: for  $i = \mu - 1$  down-to 0 do
7:    $\vec{r} \leftarrow T^{-1} \cdot \vec{r}$ ;                                            {Algorithm 9}
8:    $\vec{r} \leftarrow (-G) \cdot \vec{r}$ ;                                           {Algorithm 8}
9:    $\vec{r} += \vec{a}_i$ ;
10: end for
11: return  $R = \sum_{i=0}^{M-1} r_i X^i$ .

```

Theorem 18. *Algorithm 12 is correct, in-place and requires less than $\mathcal{O}(N \log^3 M)$ operations.*

Proof. The algorithm calls $\mu = \mathcal{O}(\frac{N}{M})$ times Algorithms 8 and 9, each call requiring less than $\mathcal{O}(M \log^3 M)$ operations by Theorems 9 and 10. \square

In the same way, we replace the blocks of Lines 6 to 9, symmetrically of Lines 23 to 26 and the block of Lines 17 to 19 of Algorithm 11, to obtain Algorithm 13. This again allows us to have some latitude in practice for different relative ratios between the degrees of the dividend and the divisor.

Theorem 19. *Algorithm 13 is correct, in-place and requires less than*

$$\mathcal{O}\left(\left(\frac{M}{m} + \log m\right) N \log^2 m\right)$$

operations.

Proof. The algorithm calls $\eta\alpha + \alpha^2 = \mathcal{O}(\frac{N}{m} \frac{M}{m})$ times Algorithm 7, each call requiring less than $\mathcal{O}(m \log^2 m)$ operations by Theorem 8. It also calls $\eta = \mathcal{O}(\frac{N}{m})$ times Algorithms 8 and 9, each call requiring less than $\mathcal{O}(M \log^3 M)$ operations by Theorems 9 and 10. \square

6 Conclusion

We have presented novel algorithms computing f-circulant and Toeplitz matrix-vector multiplications in-place. This allows us to derive novel algorithms for

Algorithm 13 In-place euclidean block remainder (fully in-place variant of Algorithm 11)

Input: $m, \vec{a} \in \mathbb{D}^N, \vec{b} \in \mathbb{D}^M, \eta = \lceil \frac{N-M+1}{m} \rceil, \alpha = \lceil \frac{M}{m} \rceil + 1;$

Output: $R = [R_0, \dots, R_{\alpha-1}]$, with $R_k = \vec{\rho}_k - \sum_{j=0}^k F_j \left(\sum_{i=1}^{\eta-j} H_i \vec{a}_{i+j} \right)$ and the H_i as defined in Theorem 14.

```

1: Let  $\eta = \lceil \frac{N-M+1}{m} \rceil, \alpha = \lceil \frac{M}{m} \rceil + 1;$ 
2: Let  $[\vec{\rho}_0, \dots, \vec{\rho}_{\alpha-1}] = [a_0, \dots, a_{M-1}, 0, \dots, 0]$  and  $[\vec{a}_1, \dots, \vec{a}_\eta] = [a_M, \dots, a_N, 0, \dots, 0]$ 
3: Let  $T = \mathcal{T}([0_{M-1}, b_M, \dots, b_1]), G = \mathcal{T}([b_{M-1}, \dots, b_0, \vec{0}_{M-1}])$ 
4: Let  $[T_1 \ G_2 \ \dots \ G_\alpha \ 0] = [I_m \ 0] [T \ G]$ 
5: for  $i = \eta$  down-to 1 do                                     {Propagate  $\eta$  times Eq. (9) to  $\vec{a}$ }
6:    $\vec{a}_i = T_1^{-1} \cdot \vec{a}_i;$                                        {Algorithm 9}
7:   for  $j = 2$  to  $\min\{\alpha; i\}$  do
8:      $\vec{a}_{i-j+1} -= G_j \cdot \vec{a}_i;$                                    {Algorithm 7}
9:   end for
10: end for
11: for  $k = 0$  to  $\alpha - 1$  do
12:   Let  $i_0 = k \cdot m;$ 
13:   Let  $[F_0 \ \dots \ F_k \ 0] = [0_{i_0} \ I_m \ 0_{M-i_0-m}] \cdot [G \ 0];$ 
14:    $R_k = \vec{\rho}_k;$ 
15:   for  $j = 0$  to  $k$  do
16:      $R_k -= F_j \cdot \vec{a}_j;$                                        {Algorithm 7}
17:   end for
18: end for
19: for  $i = 1$  to  $\eta$  do                                             {Unroll  $\eta$  times Eq. (9) to  $\vec{a}$ }
20:   for  $j = 2$  to  $\min\{\alpha; i\}$  do
21:      $\vec{a}_{i-j+1} += G_j \cdot \vec{a}_i;$                                    {Algorithm 7}
22:   end for
23:    $\vec{a}_i = T_1 \cdot \vec{a}_i;$                                        {Transpose of Algorithm 8}
24: end for
25: return  $R = [R_0, \dots, R_{\alpha-1}].$ 

```

accumulated or over-place Toeplitz multiplication or system solving. We also present algorithms that reduce the extra storage required to compute the remainder only when dividing polynomials. Eventually, we combine these techniques to propose the first in-place algorithms computing only the remainder of the polynomial euclidean division.

Further work include finding ways to (1) remove the extra logarithmic factors in the complexities; (2) deal with the field with 3 elements, there maybe some linearization of the extension with 9 elements could help; (3) avoid dividing by 2, in order to handle the even characteristic case. We propose some ways to handle most of the characteristic 2 cases in in Appendix B. Unfortunately there remains some cases we still do not have $\tilde{O}(n)$ in-place algorithms. This is \mathbb{F}_2 , \mathbb{F}_3 , and the finite fields with cyclic group cardinality a Mersenne prime (as of today, 51 are known).

References

- [1] Skander Belhaj, Marwa Dridi, and Ahmed Salam. A fast algorithm for solving banded Toeplitz systems. *Computers & Mathematics with Applications*, 70(12):2958–2967, 2015. URL: <https://www.sciencedirect.com/science/article/pii/S0898122115004939>, doi:<https://doi.org/10.1016/j.camwa.2015.10.010>.
- [2] Dario Bini and Victor Y. Pan. Fast parallel polynomial division via reduction to triangular toeplitz matrix inversion and to polynomial inversion modulo a power. *Inf. Process. Lett.*, 21(2):79–81, 1985. doi:[10.1016/0020-0190\(85\)90037-7](https://doi.org/10.1016/0020-0190(85)90037-7).
- [3] Dario Bini and Victor Y. Pan. *Polynomial and matrix computations, 1st Edition*, volume 12 of *Progress in theoretical computer science*. Birkhäuser, 1994. doi:[10.1007/978-1-4612-0265-3](https://doi.org/10.1007/978-1-4612-0265-3).
- [4] A. Bostan, G. Lecerf, and É. Schost. Tellegen’s Principle into Practice. In *Proceedings of the 2003 International Symposium on Symbolic and Algebraic Computation, ISSAC ’03*, pages 37–44, New York, NY, USA, 2003. ACM. doi:[10.1145/860854.860870](https://doi.org/10.1145/860854.860870).
- [5] Brice Boyer, Jean-Guillaume Dumas, Clément Pernet, and Wei Zhou. Memory efficient scheduling of Strassen-Winograd’s matrix multiplication algorithm. In John P. May, editor, *ISSAC’2009, Proceedings of the 2009 International Symposium on Symbolic and Algebraic Computation, Seoul, Korea*, pages 135–143, New York, July 2009. ACM Press. URL: <http://hal.archives-ouvertes.fr/hal-00163141>.
- [6] David G. Cantor and Erich Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28(7):693–701, 1991. doi:[10.1007/BF01178683](https://doi.org/10.1007/BF01178683).
- [7] Daniel Commenges and Michel Monsion. Fast inversion of triangular Toeplitz matrices. *IEEE Transactions on Automatic Control*, 29(3):250–251, 1984. doi:[10.1109/TAC.1984.1103499](https://doi.org/10.1109/TAC.1984.1103499).
- [8] Nicholas Coxon. An in-place truncated fourier transform. *Journal of Symbolic Computation*, 110:66–80, 2022. URL: <https://www.sciencedirect.com/science/article/pii/S0747717121000705>, doi:<https://doi.org/10.1016/j.jsc.2021.10.002>.
- [9] Carmine Di Fiore, Francesco Tudisco, and Paolo Zellini. Lower triangular toeplitz–ramanujan systems whose solution yields the bernoulli numbers. *Linear Algebra and its Applications*, 496:510–526, 2016. URL: <https://www.sciencedirect.com/science/article/pii/S0024379516000987>, doi:[10.1016/j.laa.2016.02.007](https://doi.org/10.1016/j.laa.2016.02.007).

- [10] Jean-Guillaume Dumas, Clément Pernet, and Alexandre Sedoglavic. Some fast algorithms multiplying a matrix by its adjoint. *Journal of Symbolic Computation*, 115:285–315, March 2023. URL: <https://hal.archives-ouvertes.fr/hal-03095393>, doi:10.1016/j.jsc.2022.08.009.
- [11] Pascal Giorgi, Bruno Grenet, and Daniel S. Roche. Generic reductions for in-place polynomial multiplication. In James H. Davenport, Dongming Wang, Manuel Kauers, and Russell J. Bradford, editors, *ISSAC'2019, Proceedings of the 2019 International Symposium on Symbolic and Algebraic Computation, Beijing, China*, pages 187–194, New York, July 2019. ACM Press. doi:10.1145/3326229.3326249.
- [12] Pascal Giorgi, Bruno Grenet, and Daniel S. Roche. Fast in-place algorithms for polynomial operations: division, evaluation, interpolation. In Ioannis Z. Emiris, Lihong Zhi, and Anton Leykin, editors, *ISSAC'2020, Proceedings of the 2020 International Symposium on Symbolic and Algebraic Computation, Kalamata, Greece*, pages 210–217, New York, July 2020. ACM Press. doi:10.1145/3373207.3404061.
- [13] Gene H. Golub and Charles F. Van Loan. *Matrix computations*. Johns Hopkins Studies in the Mathematical Sciences. The Johns Hopkins University Press, Baltimore, MD, USA, third edition, 1996.
- [14] David Harvey and Daniel S. Roche. An in-place truncated fourier transform and applications to polynomial multiplication. In Wolfram Koepf, editor, *ISSAC'2010, Proceedings of the 2010 International Symposium on Symbolic and Algebraic Computation, Munich, Germany*, page 325–329, New York, July 2010. ACM Press. doi:10.1145/1837934.1837996.
- [15] David Harvey and Joris van der Hoeven. Polynomial multiplication over finite fields in time $O(n \log n)$. *J. ACM*, 69(2):12:1–12:40, 2022. doi:10.1145/3505584.
- [16] Hsiang-Tsung Kung. On computing reciprocals of power series. *Numerische Mathematik*, 22(5):341–348, October 1974. doi:10.1007/BF01436917.
- [17] Fu-Rong Lin. An explicit formula for the inverse of band triangular Toeplitz matrix. *Linear Algebra and its Applications*, 428(2):520–534, 2008. Special Issue devoted to the Second International Conference on Structured Matrices. URL: <https://www.sciencedirect.com/science/article/pii/S0024379507002145>, doi:10.1016/j.laa.2007.05.007.
- [18] Fu-Rong Lin, Wai-Ki Ching, and Michael K. Ng. Fast inversion of triangular Toeplitz matrices. *Theor. Comput. Sci.*, 315(2-3):511–523, 2004. doi:10.1016/j.tcs.2004.01.005.

- [19] Alexander N. Malyshev and Miloud Sadkane. Fast solution of unsymmetric banded Toeplitz systems by means of spectral factorizations and woodbury’s formula. *Numerical Linear Algebra with Applications*, 21(1):13–23, 2014. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/nla.1853>, doi:10.1002/nla.1853.
- [20] Robert Moenck and Allan Borodin. Fast modular transforms via division. In *13th Annual Symposium on Switching and Automata Theory (Swat 1972)*, pages 90–96, October 1972. doi:10.1109/SWAT.1972.5.
- [21] Victor Y. Pan. *Structured Matrices and Polynomials: Unified Superfast Algorithms*. Birkhäuser, Boston, MA, USA, 2001. doi:10.1007/978-1-4612-0129-8.
- [22] Daniel S. Roche. Space-and time-efficient polynomial multiplication. In *Proceedings of the 2009 International Symposium on Symbolic and Algebraic Computation*, pages 295–302. ACM, 2009.
- [23] Andreas Rosowski. On fast computation of a circulant matrix-vector product. 2021. arXiv:2103.02605.
- [24] William F. Trench. Explicit inversion formulas for toeplitz band matrices. *SIAM Journal on Algebraic Discrete Methods*, 6(4):546–554, 1985. doi:10.1137/0606054.
- [25] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, 3rd edition, 2013.
- [26] Shmuel Winograd. On computing the discrete Fourier transform. *Proceedings of the National Academy of Sciences*, 73(4):1005–1006, 1976. doi:10.1073/pnas.73.4.1005.

A f-circulant/polynomial isomorphism

It is well known that the algebra of f -circulant $n \times n$ matrices is isomorphic to the algebra of polynomials modulo $X^n - f$, see, e.g., [21, Theorem 2.6.1]. More precisely, define the reverse of a vector of dimension n as $\text{rev}(\vec{a}) = a_{n..1}$ and identify polynomials to their vector of coefficients. Then we have that

$$P(X)Q(X) \pmod{X^n - f} = \text{rev}(\mathcal{C}_f(P) \cdot \text{rev}(Q))$$

Let $Y = X^{n/2}$, then, when f is q square, the recursive branch of Algorithm 3 is a recursive application of the Chinese Remaindering theorem with $X^n - f = (X^{n/2} + \sqrt{f})(X^{n/2} - \sqrt{f})$, or $Y^2 - f = (Y + \sqrt{f})(Y - \sqrt{f})$. Indeed both factors are coprime and satisfy

$$\frac{1}{2\sqrt{f}}(Y + \sqrt{f}) - \frac{1}{2\sqrt{f}}(Y - \sqrt{f}) = 1.$$

Therefore Let $C_+ = (C + PQ \bmod (Y + \sqrt{f}))$ and $C_- = (C + PQ \bmod (Y - \sqrt{f}))$, then $C + PQ \bmod Y^2 - f = \frac{1}{2\sqrt{f}}C_+(Y + \sqrt{f}) - \frac{1}{2\sqrt{f}}C_-(Y - \sqrt{f}) = \frac{1}{2\sqrt{f}}(C_+ - C_-)Y + \frac{1}{2}(C_+ + C_-)$.

B Characteristic two

With the isomorphism of Appendix A we can find an equivalent of Algorithm 3 for most of the finite fields of characteristic two.

Consider $q = 2^m$ and let $q - 1 = kt$, with $k \leq t$. Then consider a k -th power, f , in \mathbb{F}_q^* , that is for a generator g of this cyclic group, consider $f = g^{ik} = [ik]g$ for some i . Then the $f_j = [i + jt]g$, for $j = 1..k$ are all the k -th roots of f and they all belong to \mathbb{F}_q^* , as $(i + jt)k \equiv ik \pmod{q - 1}$ and $q - 1 = kt$.

Thus, in \mathbb{F}_q , we have that $Y^k - f = \prod_{j=1}^k (Y - f_j)$. Further, $X^{uk} - f = \prod_{j=1}^k (X^u - f_j)$ (again, let $Y = X^u$). All the $(X^u - f_j)$ are co-prime, since all the k -th roots are distinct (they have distinct exponent with respect to g , between 0 and $q - 1$).

Thus, computing the extended gcd between $Y - f_j$ and $(Y^k - f)/(Y - f_j)$, gives Bezout coefficients U_j and v_j , s.t. $U_j(Y - f_j) + v_j(Y^k - f)/(Y - f_j) = 1$, with $\deg v_j < \deg(Y - f_j) = 1$. This means that $v_j \in \mathbb{F}_q^*$ (and this remains true if $Y = X^u$, for any u).

Let $n = ku$, the Chinese remaindering Lagrange formula now writes:

$$c \bmod (X^n - f) = \prod_{j=1}^k (c \bmod (X^u - f_j)) v_j \frac{X^n - f}{X^u - f_j}.$$

Thus, as soon as the number of invertibles in \mathbb{F}_q , $2^m - 1$ is not a Mersenne prime, then it can be written as $2^m - 1 = kt$ with $t \geq k > 1$ (that is $t \geq k \geq 3$) and we can form a k -adic recursion for circulant matrices. This gives a fast algorithm: if $T(n) \leq kT(n/k) + O(n \log^\alpha(n))$, then $T(n) = \mathcal{O}(n \log^{\alpha+1}(n))$.

Further, if $k > 1$, then 1 is not the only k -th power (there are $t \geq 3$ of them), and we can built the analogous of the full Algorithm 3. Indeed, by letting $\lambda = (b^k - f)/(f - 1)$, for another k -th power b^k , we can always compute $\mathcal{C}_f = \frac{\lambda}{\lambda+1} \mathcal{C}_1 + \frac{1}{\lambda+1} \mathcal{C}_{b^k}$, and the accumulation $c = c + \mathcal{C}_f$ can be computed by

$$c = \frac{\lambda}{\lambda+1} \left(\frac{1}{\lambda} ((\lambda+1)c + \mathcal{C}_{b^k}) + \mathcal{C}_1 \right).$$

The conditions are that λ exists, thus $f \neq 1$, λ invertible, thus f is not a k -th power, $\lambda + 1$ invertible, thus $b^k \neq 1$.

We have proven:

Proposition 20. *For $A(X)$ and $B(X)$ polynomials of respective degrees N and M over any field with even characteristic $q = 2^m$, such that $2^m - 1$ is not prime, there exist an in-place algorithm requiring less than $\mathcal{O}(N \log^3 M)$ arithmetic operations.*

C Open cases

Cyclic group of Mersenne prime cardinality In case of a Mersenne prime, then unfortunately the only power with all its roots in the field is 1 and we thus cannot combine recursive calls. This prevents us for now to have superfast algorithms in $\mathbb{F}_2, \mathbb{F}_4, \mathbb{F}_8, \mathbb{F}_{32}, \mathbb{F}_{128}, \mathbb{F}_{8192}, \dots$, (as of today, only 51 Mersenne primes are known¹, and we do not know if there are an infinite number of those or not.

Still, in characteristic 2, one can always use a triadic algorithm in an extension of degree 2, that is with an algebraic root of $j^2 + j + 1$. Then the 3 roots of unity, 1, j and $j + 1$ are by definition inside the extension. And all the coefficients can be formed as polynomials of degree 1 in j .

Now, to remain in-place, one cannot double the space of each coefficient (that would be required to store polynomials of degree 1). But it should be possible to sequentially compute with the real and imaginary parts. In fact, the algorithm of [26, Eq. (19)] does exactly that, i.e. computing modulo $Y - 1, Y - j$ and $Y - (j + 1)$ and then explicitly reconstructing the polynomial of degree 2 in $Y = X^{n/3}$. But this algorithm uses 4 multiplications, that is 4 recursive calls, with polynomials of degree $n/3$. Therefore its complexity bound would only be $\mathcal{O}(n^{\log_3(4)})$.

Again our approach to have an algorithm for each possible coefficients in the field, then works if there exist at least 2 distinct cubes in the extension. This just discards \mathbb{F}_2 and F_4 .

Over \mathbb{F}_3 The same approach should work for -1 over \mathbb{F}_3 , as it becomes a square in $\mathbb{F}_9 \cong \mathbb{F}_3[j]/(j^2 + 1)$. But, since Gauß, the minimum number of multiplications required to multiply two complex numbers is 3, resulting in the Karatsuba $\mathcal{O}(n^{\log_2(3)})$ complexity.

¹<https://www.mersenne.org/primes>