

SPACE-EFFICIENT QUANTUM MULTIPLICATION OF POLYNOMIALS FOR BINARY FINITE FIELDS WITH SUB-QUADRATIC TOFFOLI GATE COUNT

IGGY VAN HOOF

*Mathematics and Computer Science, Technische Universiteit Eindhoven
Groene Loper 5, 5612 AE, Eindhoven, Netherlands*

Received April 9, 2020

Revised July 20, 2020

Multiplication is an essential step in a lot of calculations. In this paper we look at multiplication of 2 binary polynomials of degree at most $n - 1$, modulo an irreducible polynomial of degree n with $2n$ input and n output qubits, without ancillary qubits, assuming no errors. With straightforward schoolbook methods this would result in a quadratic number of Toffoli gates and a linear number of CNOT gates. This paper introduces a new algorithm that uses the same space, but by utilizing space-efficient variants of Karatsuba multiplication methods it requires only $O(n^{\log_2(3)})$ Toffoli gates at the cost of a higher CNOT gate count: theoretically up to $O(n^2)$ but in examples the CNOT gate count looks a lot better.

Keywords: Quantum algorithm, Karatsuba multiplication, Binary polynomials

Communicated by: R Jozsa & M Mosca

1 Introduction

Multiplication of two polynomials in a finite field is an important step in many algorithms, such as point addition in elliptic curve cryptography. For classical computers a wealth of variations exist, often based around Karatsuba’s multiplication method [1].

In the classical setting, temporary results for the steps of Karatsuba calculations have traditionally been stored separately. In 1993 Maeder [2] used around $2n$ additional space for multiplying degree- n polynomials. This was improved by Thomé in 2002 to n temporary space which at the time was believed to be optimal: “it does not seem likely that anything better than this result can be obtained.” [3] However, in 2009 Roche did obtain a better result: $O(\log n)$ space Karatsuba multiplication of polynomials without additional time by doing many in-place operations [4]. This was expanded by Cheng [5] to also work for integers. Despite the advantages these variants offer, these methods are still relatively unknown.

This bound of $O(\log n)$ temporary storage is still higher than the bound presented in this paper, which is reduced to 0 by partly overwriting the input polynomial and restoring it before the end. With this advantage we can modify the algorithms presented by Roche [4] for the quantum setting. The algorithms in this paper have an exponential speedup over other quantum algorithms that do not use extra space [6]. Other variants that reach the same speedup as classical Karatsuba multiplication in the quantum setting so far have done

so at the cost of space [7].

1.1 Overview

We introduce our notation for quantum computing by giving the elementary quantum gates in section 2. Our new multiplication algorithm needs several subroutines, specifically modular shifts and multiplication by a constant polynomial, introduced in section 3. We introduce a Quantum Karatsuba algorithm for multiplication without reduction in section 4 and in binary finite fields in section 5. Both algorithms run without ancillary qubits and have a sub-quadratic Toffoli gate count. We implemented the algorithm in a simulated quantum computer and present the gate counts for specific finite fields in section 6.

2 Quantum background

Quantum computing uses reversible gates, which unlike classical gates can be run in reverse and require an equal number of input and output quantum bits (qubits). In this paper we will not make use of the quantum properties of qubits, but the gates we use can be applied to superpositions of qubits in states 1 and 0. For the purpose of multiplication we need two gates to do reversible addition and multiplication:

- The CNOT, or Feynman, gate serves as the equivalent of XOR or \mathbb{F}_2 -addition. This gate takes 2 qubits as inputs and adds one input to the other qubit and outputs the other qubit as itself: $(a, b) \rightarrow (a, a \oplus b)$. It is reversible and its own inverse: applying it twice would result in $(a, a \oplus b \oplus a) = (a, b)$. In Figure 1 an example has been drawn. In algorithms we write this as $\text{CNOT}(a, b)$.

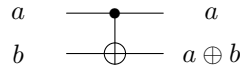


Fig. 1. The CNOT gate

- The Toffoli (TOF) gate serves as the equivalent of AND or \mathbb{F}_2 -multiplication in our case. This gate takes 3 qubits as inputs and adds the result of multiplication of the first two qubits to the third qubit and outputs the other qubits as themselves: $(a, b, c) \rightarrow (a, b, c \oplus (a \cdot b))$. It is also its own inverse. In Figure 2 an example has been drawn. In algorithms we write this as $\text{TOF}(a, b, c)$

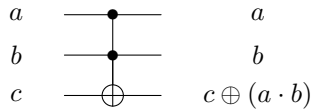


Fig. 2. The TOF gate

In addition to these operations, we will also need to swap some qubits. Unlike the previous gates we do not build these in physical circuits. Rather, we change the index on some qubits: if we were to swap qubits 1 and 2 we would simply refer to qubit 1 as “2” and qubit 2 as “1” from that point on without counting any quantum gates. In Figure 3 an example has been



Fig. 3. The swap

drawn.

These 3 actions are the only essential ones we use in this paper. Although none of these are explicit quantum actions, the quantum dimension comes from optimizing for low Toffoli gate count. Currently no large quantum computer exists but current estimates put the cost of one Toffoli gate at many times that of a CNOT gate.

3 Basic Arithmetic

In this section we discuss reversible in-place algorithms for the basic arithmetic of binary polynomials.

3.1 Addition and binary shift

The first operation we consider, addition, can easily be implemented for binary polynomials. Individual additions can be done with a CNOT gate, the addition of two polynomials of degree at most n takes $n + 1$ CNOT gates with depth 1. This operation uses ancillary qubits and the result of the addition replaces either of the inputs. Since addition is component-wise, addition for polynomials over \mathbb{F}_2 is the same as addition for elements of the field \mathbb{F}_{2^n} .

Binary shifts are straightforward: they correspond to multiplying or dividing by x . This requires no quantum computation by doing a series of swaps.

Finally, if we have a fixed n , a polynomial $g(x)$ of degree at most $n - 1$ and want to do a multiplication by x followed by a modular reduction by a fixed weight- ω and degree- n polynomial $m(x)$ that has coefficient 1 for x^0 , we can do this in 2 steps. We represent $m(x)$ as M where M is an ordered list of length ω that contains the degrees of the nonzero terms in descending order, for example if $m(x) = 1 + x^3 + x^{10}$ we get $M = [10, 3, 0]$. Let $g(x) = \sum_{i=0}^{n-1} g_i x^i$:

- Step 1: For every qubit g_i change its index so that it represents the coefficient of $x^{i+1 \bmod n}$. Let h_i be the coefficients of the relabeled polynomial, i.e. $h_{i+1 \bmod n} = g_i$.

Algorithm 1: MODSHIFT $_{m(x)}$ Reversible algorithm for in-place multiplication by x in $\mathbb{F}_2[x]/m(x)$ with $m(x)$ an irreducible binary polynomial.

Fixed input : A constant irreducible binary polynomial $m(x)$ of degree n , represented as an ordered list M of length ω that contains the degrees of the nonzero terms in descending order.

Quantum input: A binary polynomial $g(x)$ of degree up to $n - 1$ stored in an array G .

Result: G as $g \cdot x$ in the field $\mathbb{F}_2/m(x)$.

```

1 for  $i = 1..n$  do
2    $\lfloor$  SWAP( $G[n - i], G[n - i - 1]$ )
3 for  $i = 1..\omega - 2$  do
4    $\lfloor$  CNOT( $G[0], G[M_i]$ )
```

- Step 2: Apply CNOT controlled by the x^0 term h_0 (g_{n-1} before Step 1) to h_j , with $j = M_1, \dots, M_{\omega-2}$. In the example of $1 + x^3 + x^{10}$ we would apply 1 CNOT to h_3 controlled by h_0 .

See Algorithm 1 for pseudocode and Figure 4 for an example. After a multiplication by x the coefficient of x^0 is always 0. Since $m(x)$ always has coefficient 1 for x^0 , after a reduction by $m(x)$ that qubit will be 1 and if no reduction takes place that qubit is 0, which means our modular shift algorithm is always reversible. This results in a total of $\omega - 2$ CNOT gates for a modular reduction, with depth $\omega - 2$ and we do not use ancillary qubits. Since we use reversible gates, running this circuit in reverse corresponds to dividing by x modulo $m(x)$.

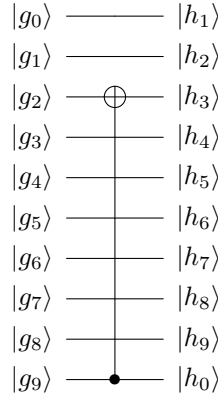


Fig. 4. Binary shift circuit for $\mathbb{F}_{2^{10}}$ with $g_0 + \dots + g_9x^9$ as the input and $h_0 + \dots + h_9x^9 = g_9 + g_0x + g_1x^2 + (g_2 + g_9)x^3 + g_3x^4 + \dots + g_9x^9$ as the output.

3.2 Multiplication by a constant polynomial

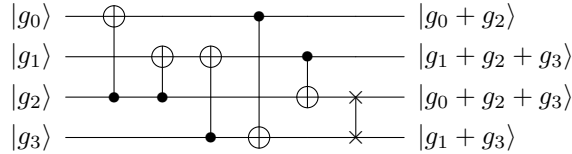


Fig. 5. Multiplication of g by $1 + x^2$ modulo $1 + x + x^4$. Depth 4 and 5 CNOT gates.

Multiplication by a constant non-zero polynomial in a fixed binary field is \mathbb{F}_2 -linear: as the field polynomial is irreducible, every input corresponds to exactly one output. We can see that any such multiplication can be represented as a matrix, which we can turn into a circuit using an *LUP*-decomposition, an algorithm also used by Amento, Rötteler and Steinwandt [8]. For example, multiplication by $1 + x^2$ modulo $1 + x + x^4$ can be represented by a matrix Γ . Using the decomposition $\Gamma = P^{-1}LU$ we get an upper and lower triangular matrix which we can translate into a circuit. Any 1 not on the diagonal in U and L is a CNOT controlled by the column number on the row number. In cases of conflict^a for U CNOT gates should be

^aConflicts exist if according to the triangular matrix a CNOT would both have to be applied on and controlled by a qubit. By doing the controlled operation first and applying the operation on it afterwards, we ensure that the matrix multiplication is correctly translated.

performed top row first, second row second and so on and for L CNOT gates from the bottom row up. P represents a series of swaps, and can be represented either as a permutation matrix or an ordered list with all elements from 0 to $n - 1$.

$$\Gamma = \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \end{pmatrix} = P^{-1}LU = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (1)$$

Figure 5 shows how we translate Γ . According to [8] this costs up to $n^2 + n$ CNOT gates with depth up to $2n$. We can improve this count by noting L and U are each size n by n and can have up to $(n^2 - n)/2$ non-diagonal non-zero entries, giving us up to $n^2 - n$ CNOT gates. Note that the LUP -decomposition is precomputed and for any fixed polynomial and field we can give an exact CNOT gate count and depth.

Since this algorithm is introduced in [8] without correctness proof and we will use it later for a bigger algorithm, we will write an explicit implementation and go over the correctness of this algorithm. Note that since we are working with reversible algorithms, multiplying by constant $f(x)$ is the same as doing the reverse of multiplying by constant $f(x)^{-1}$.

Theorem 1: Algorithm 2 correctly describes multiplication by a non-zero constant polynomial in a fixed binary field.

Proof: Since multiplication by a non-zero constant in a finite field is a linear map, an invertible matrix Γ to represent this linear map must exist. Since Γ is invertible, its decomposition L, U, P^{-1} must also consist of linear maps. Since we are working in a binary field and U is an invertible upper-triangular matrix, the diagonal of U is all-one. If we look at lines 1 through 4 of the algorithm, we can see it corresponds to applying linear map U to g , as it results in $g_i = \sum_{j=0}^{n-1} u_{i,j}g_j$ for $i = 0, \dots, n - 1$. Analogously the same is true for L in lines 5 through 8. We can also see that if P^{-1} is a row-permutation of the identity matrix, lines 9 through 13 will apply it correctly. Since $P^{-1}LU = \Gamma$ we have correctly applied the linear map Γ . \square

Note that the algorithm is not optimized for depth, for example in circuit 5 the first and second CNOT could be swapped so the depth would be 3 rather than 4.

3.2.1 Choice of field polynomials

When doing operations in a finite binary field we can choose what representation we use, as long as the polynomial $m(x)$ is irreducible. Our goal is to make the matrices L and U as sparse as possible. For this purpose we also want our Γ to be as sparse as possible, which can be achieved in two steps: pick irreducible polynomials with as few non-zero coefficients as possible, i.e. trinomials when available and pentanomials otherwise, and pick irreducible polynomials where the second highest non-constant term has the lowest possible degree. For example, the pentanomial $1 + x^3 + x^4 + x^{19} + x^{20}$ would require 108 CNOT gates, the pentanomial $1 + x^3 + x^5 + x^9 + x^{20}$ would require 55 CNOT gates, while the trinomial $1 + x^3 + x^{20}$ would require only 27. All 3 polynomials are irreducible. In Table 1 we can see some examples of gate counts for various choices of n . The depth count is an upper bound without accounting for swapping gates.

Algorithm 2: CONSTMODMULT $_{f(x),m(x)}$, from [8]. Reversible algorithm for in-place multiplication by a nonzero constant polynomial $f(x)$ in $\mathbb{F}_2[x]/m(x)$ with $m(x)$ an irreducible polynomial.

Fixed input : A binary LUP -decomposition L, U, P^{-1} for a binary n by n matrix that corresponds to multiplication by the constant polynomial $f(x)$ in the field $\mathbb{F}_2[x]/m(x)$.

Quantum input: A binary polynomial $g(x)$ of degree up to $n - 1$ stored in an array G .

Result: G as $f \cdot g$ in the field $\mathbb{F}_2/m(x)$.

```

1 for  $i = 0..n - 1$  //  $U \cdot G$ 
2 do
3   for  $j = i + 1..n - 1$  do
4     if  $U[i, j] = 1$  then
5       CNOT( $G[j], G[i]$ )
6 for  $i = n - 1..0$  //  $L \cdot UG$ 
7 do
8   for  $j = i - 1..0$  do
9     if  $L[i, j] = 1$  then
10      CNOT( $G[j], G[i]$ )
11 for  $i = 0..n$  //  $P^{-1} \cdot LUG$ 
12 do
13   for  $j = i + 1..n - 1$  do
14     if  $P^{-1}[i, j] = 1$  then
15       SWAP( $G[i], G[j]$ )
16       SWAP column  $i$  and  $j$  of  $P^{-1}$ 

```

4 Quantum Multiplication for binary polynomials

This section details schoolbook multiplication and we present our new Karatsuba algorithm.

4.1 Quantum Schoolbook Multiplication

The simplest way to multiply is schoolbook multiplication. For two polynomials of degree at most $n - 1$ that takes n^2 Toffoli gates, the number of pairs of qubits from the first and second polynomial. While the computation does not use ancillary qubits, the result needs to be stored separately from input in $2n - 1$ qubits; unlike the previous circuits we cannot replace either of the inputs with the result since the Toffoli gate requires a separate output. If we want to apply modular reduction steps by a weight- k and degree- n odd polynomial, this adds $(n - 1) \cdot (k - 2)$ CNOT gates and uses no ancillary qubits (by using the modular shift algorithm after every n multiplications). The result is stored in n qubits.

4.2 Classic Karatsuba multiplication in binary polynomial rings

Rather than using schoolbook multiplication, methods like Karatsuba multiplication [1] can speed up multiplication of large numbers. We can look at in-place multiplication in the

Degree	Irreducible polynomial	Source	CNOT gates	Depth upper bound
4	[4, 1, 0]	[9]	5	4
8	[8, 4, 3, 1, 0]	[9]	20	14
16	[16, 5, 3, 1, 0]	[9]	47	30
32	[32, 7, 3, 2, 0]	[9]	133	93
64	[64, 4, 3, 1, 0]	[9]	264	182
127	[127, 1, 0]	[9]	396	293
128	[128, 7, 2, 1, 0]	[9]	626	443
163	[163, 7, 6, 3, 0]	[10]	740	975
163	[163, 89, 74, 15, 0]	[11]	1885	1646
233	[233, 74, 0]	[10]	3319	2976
256	[256, 10, 5, 2, 0]	[9]	1401	1030
283	[283, 12, 7, 5, 0]	[10]	2117	1700
283	[283, 160, 123, 37, 0]	[11]	6785	6368
571	[571, 10, 5, 2, 0]	[10]	4027	3177
571	[571, 353, 218, 135, 0]	[11]	33182	32331
1024	[1024, 19, 6, 1, 0]	[12]	8147	6624

Table 1. Comparison of the CNOT gates required for various instances of Algorithm 2. Source is the source of the polynomial.

classical case for ideas [4]. As input we take two polynomials of size up to n , $f(x)$ and $g(x)$ as well as a polynomial of size $2n$: $h(x)$. As output we desire $h + f \cdot g$. For some k such that $\frac{n}{2} \leq k < n$ (we will always use $k = \lceil \frac{n}{2} \rceil$) we can split each polynomial as follows: $f = f_0 + f_1x^k$, $g = g_0 + g_1x^k$ and $h = h_0 + h_1x^k + h_2x^{2k} + h_3x^{3k}$.

We compute intermediate products $\alpha = f_0 \cdot g_0$, $\beta = f_1 \cdot g_1$ and $\gamma = (f_0 + f_1) \cdot (g_0 + g_1)$. Finally, we add these in the right way for Karatsuba multiplication:

$$h + f \cdot g = h + \alpha + (\gamma + \alpha + \beta)x^k + \beta x^{2k} \quad (2)$$

For cleanliness, we can split up our α, β, γ in the same way as f and g to get a result with no overlap, which is useful for checking correctness:

$$h + f \cdot g = (h_0 + \alpha_0) + (h_1 + \alpha_0 + \alpha_1 + \beta_0 + \gamma_0)x^k + (h_2 + \alpha_1 + \beta_0 + \beta_1 + \gamma_1)x^{2k} + (h_3 + \beta_1)x^{3k} \quad (3)$$

Alternatively, we can rewrite this another way that will prove useful:

$$h + f \cdot g = h + (1 + x^k)\alpha + x^k\gamma + x^k(1 + x^k)\beta. \quad (4)$$

4.3 Reversible Karatsuba multiplication in binary polynomial rings

Based on these equations we can split our multiplication algorithm into 2 parts: given $f(x), g(x), h(x)$ calculate $h + f \cdot g$ and given $k, f(x), g(x), h(x)$ with $k > \max(\deg(f), \deg(g))$ calculate $h + (1 + x^k)f \cdot g$. We will look at our algorithms for the 2 parts, which can then be used recursively to provide a significant improvement to the schoolbook algorithm in terms of Toffoli gate count.

Lemma 1: Let $k = \lceil \frac{n}{2} \rceil$. Given polynomials f, g of degree up to $n - 1$ with $n > 1$ and h of degree up to $2n - 2$. Assuming Algorithm 4 correctly calculates $h' + (1 + x^k)f'g'$ for f', g' up

Algorithm 3: KMULT_n . Reversible algorithm for multiplication of 2 polynomials of degree up to $n - 1$.

Fixed input : A constant integer n to indicate polynomial size and an integer $k < n \leq 2k$ with $k = \lceil \frac{n}{2} \rceil$ for $n > 1$ and $k = 0$ for $n = 1$, to indicate upper and lower half.

Quantum input: Two binary polynomial f, g of degree up to $n - 1$ stored in arrays F and G respectively of size n . A binary polynomial h of degree up to $2n - 2$ stored in array H of size $2n - 1$.

Result: F and G as input, H as $h + fg$

```

1 if  $n > 1$  then
2    $\text{KMULT1xk}_{k,k}(F[0..k-1], G[0..k-1], H[0..3k-2])$ 
3    $\text{KMULT1xk}_{k,n-k}(F[k..n-1], G[k..n-1], H[k..2n-2])$ 
4    $\text{CNOT}(F[k..n-1], F[0..n-k-1])$ 
5    $\text{CNOT}(G[k..n-1], G[0..n-k-1])$ 
6    $\text{KMULT}_k(F[0..k-1], G[0..k-1], H[k..3k-2])$ 
7    $\text{CNOT}(G[k..n-1], G[0..n-k-1])$ 
8    $\text{CNOT}(F[k..n-1], F[0..n-k-1])$ 
9 else
10   $\text{TOF}(F[0], G[0], H[0])$ 

```

to degree $k - 1$ and h' up to degree $3k - 2$, and Algorithm 3 correctly calculates $h'' + f''g''$ with f'', g'' of degree $k - 1$ and h'' of degree $2k - 2$ without altering the values of f'' and g'' . Then Algorithm 3 correctly calculates $h + fg$ in $\mathbb{F}_2[x]$. The values of f and g are the same after the algorithm as they were before.

Proof: Table 2 gives the result of each line on array H , split into 4 parts of size k, k, k and $2n - 1 - 3k$ respectively: $h = h_0 + h_1x^k + h_2x^{2k} + h_3x^{3k}$. As can be seen in the table, the final result corresponds to $(h_0 + \alpha_0) + (h_1 + \alpha_0 + \alpha_1 + \beta_0 + \gamma_0)x^k + (h_2 + \alpha_1 + \beta_0 + \beta_1 + \gamma_1)x^{2k} + (h_3 + \beta_1)x^{3k} = h + f \cdot g$ as discussed in Section 4.2. Lines 7 and 8 are the inverses of lines 4 and 5 so return F and G to their original states. \square

Line	H in KMULT_n			
	$H[0..k-1]$	$H[k..2k-1]$	$H[2k..3k-1]$	$H[3k..2n-2]$
1	h_0	h_1	h_2	h_3
2	$h_0 + \alpha_0$	$h_1 + \alpha_0 + \alpha_1$	$h_2 + \alpha_1$	h_3
3-5	$h_0 + \alpha_0$	$h_1 + \alpha_0 + \alpha_1 + \beta_0$	$h_2 + \alpha_1 + \beta_0 + \beta_1$	$h_3 + \beta_1$
6-8	$h_0 + \alpha_0$	$h_1 + \alpha_0 + \alpha_1 + \beta_0 + \gamma_0$	$h_2 + \alpha_1 + \beta_0 + \beta_1 + \gamma_1$	$h_3 + \beta_1$

Table 2. Step by step calculation of Algorithm 3.

Algorithm 3 computes $h + fg$ with $4(n - k)$ CNOT gates, at a depth of 4, 1 call to itself for a k -by- k multiplication, 1 call to Algorithm 4 for a k -by- k multiplication and 1 call to Algorithm 4 for an $(n - k)$ -by- $(n - k)$ multiplication. For $n = 1$ we just have a single TOF gate.

Algorithm 4: KMULT1xk_{k,n}. Reversible algorithm for multiplication of the product of 2 polynomials of degree up to $n - 1$ by the polynomial $1 + x^k$.

Fixed input : A constant integer $k > 0$ to indicate part size as well as an integer $n \leq k$ to indicate polynomial size. $\ell = \max(0, 2n - 1 - k)$ is the size of h_2 and $(fg)_1$. In the case of Karatsuba we will have either $n = k$ or $n = k - 1$.

Quantum input: Two binary polynomials $f(x), g(x)$ of degree up to $n - 1$ stored in arrays F and G respectively of size n . A binary polynomial $h(x)$ of degree up to $k + 2n - 2$ stored in array H of size $2k + \ell$.

Result: F and G as input, H as $h + (1 + x^k)fg$

```

1 if  $n > 1$  then
2   CNOT( $H[2k..2k + \ell - 1], H[k..k + \ell - 1]$ )
3   CNOT( $H[k..2k - 1], H[0..k - 1]$ )
4   KMULTn( $F[0..n - 1], G[0..n - 1], H[k..2k + \ell - 1]$ )
5   CNOT( $H[k..2k - 1], H[0..k - 1]$ )
6   CNOT( $H[2k..2k + \ell - 1], H[k..k + \ell - 1]$ )
7 else
8   CNOT( $H[k], H[0]$ )
9   TOF( $F[0], G[0], H[k]$ )
10  CNOT( $H[k], H[0]$ )

```

Lemma 2: Given polynomials f, g of degree up to $n - 1$ with $n > 1$, polynomial h of degree up to $k + 2n - 2$ with some $k \geq n$ and assuming Algorithm 3 correctly calculates $h + fg$ with degrees of f, g and h bounded as above, Algorithm 4 correctly calculates $h + (1 + x^k)fg$ in $\mathbb{F}_2[x]$ without altering the values of f and g .

Proof: Let $\ell = \max(0, 2n - 1 - k)$. Table 3 gives the result of each step on array H , split into 3 parts of size k, k and $\ell - 1$ respectively: $h = h_0 + h_1x^k + h_2x^{2k}$. The final result corresponds to $h_0 + (fg)_0 + (h_1 + (fg)_0 + (fg)_1)x^k + (h_2 + (fg)_1)x^{2k} = h_0 + h_1x^k + h_2x^{2k} + fg + fgx^k = h + (1 + x^k)fg$, where $(fg)_0$ is the first k terms of $f \cdot g$ and $(fg)_1$ is the last up to ℓ terms.

Line	H in KMULT1xk _{k,n}		
	$H[0..k - 1]$	$H[k..2k - 1]$	$H[2k..2k + \ell - 1]$
1	h_0	h_1	h_2
2	h_0	$h_1 + h_2$	h_2
3	$h_0 + h_1 + h_2$	$h_1 + h_2$	h_2
4	$h_0 + h_1 + h_2$	$h_1 + h_2 + (fg)_0$	$h_2 + (fg)_1$
5	$h_0 + (fg)_0$	$h_1 + h_2 + (fg)_0$	$h_2 + (fg)_1$
6	$h_0 + (fg)_0$	$h_1 + (fg)_0 + (fg)_1$	$h_2 + (fg)_1$

Table 3. Step by step calculation of Algorithm 4.

f and g do not have their values altered because arrays F and G remain unchanged. \square

Algorithm 4 computes $h + (1 + x^k)fg$ with at most $2k + 2\ell \geq 2k + 2(2n - 1 - k) = 4n - 2$ CNOT gates, at a depth of 4 per layer and 1 call to Algorithm 3 for an n -by- n multiplication. For $n = 1$ both the depth and number of gates is 2 CNOT and 1 TOF gates.

Theorem 2: Given polynomials f, g of degree up to $n - 1$ and h of degree up to $2n - 2$, Algorithm 3 correctly calculates $h + fg$. The values of f and g are the same after the algorithm as they were before.

Proof: We use proof by induction. For $n = 1$ line 10 of Algorithm 3 correctly calculates $h + fg$ without altering f or g .

For $n = 2$ two calls are made to Algorithm 4 and one call to Algorithm 3 with $n' = 1$ and $k' = 1$. Lines 7-9 of Algorithm 4 correctly calculate $h' + (1 + x^k)f'g'$.

For $n > 2$ we use lemmas 1 and 2 as our inductive steps. Every time Algorithm 3 is called recursively to calculate $h' + f'g'$ with f', g' of degree $n' - 1$, it is with either $n' = \lceil \frac{n}{2} \rceil$ or $n' = n - \lceil \frac{n}{2} \rceil = \lfloor \frac{n}{2} \rfloor$.

The series $\lceil \frac{n}{2} \rceil, \lceil \frac{\lceil \frac{n}{2} \rceil}{2} \rceil, \lceil \frac{\lceil \frac{\lceil \frac{n}{2} \rceil}{2} \rceil}{2} \rceil, \dots$ reaches 1 in $O(\log n)$ steps and $\lfloor \frac{n}{2} \rfloor \leq \lceil \frac{n}{2} \rceil$. From this we can see that we reach $n' = 1$ or 2 in finite steps. By induction Algorithm 3 correctly calculates $h + fg$ and returns f and g to their original values. \square

5 Reversible Karatsuba multiplication in binary finite fields

With this basis, we can move on to the modular multiplication. We will need Algorithm 2, which we will also run in reverse for multiplication by an inverse, and the binary shifts from Section 3.1, which we will refer to as MODSHIFT, as well as the previous Karatsuba algorithms. Unlike before, we will assume we start with an all-zero input. We can see in Algorithm 11 the number of operations we use:

- 3 calls to Algorithm 3: twice for k -by- k multiplication and once for $(n - k)$ -by- $(n - k)$ multiplication.
- 2 calls to Algorithm 2 (once in reverse), each time for multiplication by the same polynomial $1 + x^k$.
- k calls to Algorithm 1 for multiplication by x^k .
- 4 times $(n - k)$ CNOT gates, half of which can be performed at the same time.

Note that Algorithm 3 can multiply two polynomials f and g of degree at most $\lceil \frac{n}{2} \rceil - 1$ while needing n space for the output polynomial h , which has degree $n - 1$ at most in the case that n is odd. We make recursive calls to Algorithm 3 rather than Algorithm 11 because it uses significantly fewer CNOT operations and fits in the required space.

Theorem 3: Algorithm 11 correctly calculates fg in a field $\mathbb{F}_2[x]/m(x)$ and the values of f and g are the same after the algorithm as they were before.

Proof: Table 4 gives the result of each line on array H . As can be seen in the table, the final result corresponds to $(1 + x^k)\alpha + x^k\gamma + x^k(1 + x^k)\beta \pmod{m}$. Lines 4 and 5 are the inverses

Algorithm 5: MODMULTZERO $_{m(x)}$. Reversible algorithm for multiplication of 2 polynomials in $\mathbb{F}_2[x]/m(x)$ with $m(x)$ an irreducible polynomial.

Fixed input : A constant integer n to indicate field size, $k = \lceil \frac{n}{2} \rceil$. $m(x)$ of degree n as the field polynomial. The LUP -decomposition precomputed for multiplication by $1 + x^k$ modulo $m(x)$.

Quantum input: Two binary polynomials $f(x), g(x)$ of degree up to $n - 1$ stored in arrays F and G respectively of size n . An all-zero array H of size n

Result: F and G as input, H as $f \cdot g \bmod m$.

```

1 CNOT( $F[k..n - 1]$ ,  $F[0..n - k - 1]$ )
2 CNOT( $G[k..n - 1]$ ,  $G[0..n - k - 1]$ )
3 KMULT $_k$ ( $F[0..k - 1]$ ,  $G[0..k - 1]$ ,  $H[0..n - 1]$ )
4 CNOT( $G[k..n - 1]$ ,  $G[0..n - k - 1]$ )
5 CNOT( $F[k..n - 1]$ ,  $F[0..n - k - 1]$ )
6 CONSTMODMULT $_{1+x^k, m(x)}^{-1}$ ( $H[0..n - 1]$ )
7 KMULT $_{n-k}$ ( $F[k..n - 1]$ ,  $G[k..n - 1]$ ,  $H[0..n - 1]$ )
8 for  $i = 0..k - 1$  do
9    $\lfloor$  MODSHIFT $_{m(x)}$ ( $H[0..n - 1]$ )
10 KMULT $_k$ ( $F[0..k - 1]$ ,  $G[0..k - 1]$ ,  $H[0..n - 1]$ )
11 CONSTMODMULT $_{1+x^k, m(x)}$ ( $H[0..n - 1]$ )

```

Line	H in MODMULTZERO $_{m(x)}$
1,2	0
3-5	γ
6	$(1 + x^k)^{-1} \gamma \bmod m$
7	$(1 + x^k)^{-1} \gamma + \beta \bmod m$
9	$x^k((1 + x^k)^{-1} \gamma + \beta) \bmod m$
10	$\alpha + x^k((1 + x^k)^{-1} \gamma + \beta) \bmod m$
11	$(1 + x^k) \alpha + x^k \gamma + x^k(1 + x^k) \beta \bmod m$

Table 4. Step-by-step calculation of Algorithm 11.

of lines 1 and 2 so return F and G to their original states. \square

Algorithm 6 provides a version of Algorithm 11 without the requirement of an all-zero H , with Table 5 giving the results of each line. Due to the k added calls to MODSHIFT it uses k more CNOT gates if $m(x)$ is a trinomial and $3k$ more CNOT gates if $m(x)$ is a pentanomial compared to Algorithm .

6 Results

Algorithm 11 uses the same number of Toffoli gates as regular Karatsuba multiplication: 3 half-sized multiplications. This means the asymptotic number of Toffoli gates is the same as for regular Karatsuba: $O(n^{\log(3)}) \approx O(n^{1.58})$. This is a significant improvement over the n^2 Toffoli gates required for schoolbook multiplication. The number of CNOT gates is less clear as the number of CNOT gates required for the multiplications with constant polynomials is strongly dependent on our choice of field polynomial. It is not within the scope of this paper to find a stronger bound than $O(n^2)$ for the number of CNOT gates, which is currently used

Algorithm 6: $\text{MODMULT}_{m(x)}$. Reversible algorithm for multiplication of 2 polynomials in $\mathbb{F}_2[x]/m(x)$ with $m(x)$ an irreducible polynomial and adding the product to a third polynomial.

Fixed input : A constant integer n to indicate field size, $k = \lceil \frac{n}{2} \rceil$. $m(x)$ of degree n as the field polynomial. The LUP -decomposition precomputed for multiplication by $1 + x^k$ modulo $m(x)$.

Quantum input: Three binary polynomials $f(x), g(x), h(x)$ of degree up to $n - 1$ stored in arrays F, G and H respectively of size n .

Result: F and G as input, H as $h + f \cdot g \mod m$.

```

1 for  $i = 0..k - 1$  do
2    $\lfloor \text{MODSHIFT}_{m(x)}^{-1}(H[0..n - 1])$ 
3    $\text{CNOT}(F[k..n - 1], F[0..n - k - 1])$ 
4    $\text{CNOT}(G[k..n - 1], G[0..n - k - 1])$ 
5    $\text{KMULT}_k(F[0..k - 1], G[0..k - 1], H[0..n - 1])$ 
6    $\text{CNOT}(G[k..n - 1], G[0..n - k - 1])$ 
7    $\text{CNOT}(F[k..n - 1], F[0..n - k - 1])$ 
8    $\text{CONSTMODMULT}_{1+x^k, m(x)}^{-1}(H[0..n - 1])$ 
9    $\text{KMULT}_{n-k}(F[k..n - 1], G[k..n - 1], H[0..n - 1])$ 
10  for  $i = 0..k - 1$  do
11     $\lfloor \text{MODSHIFT}_{m(x)}(H[0..n - 1])$ 
12     $\text{KMULT}_k(F[0..k - 1], G[0..k - 1], H[0..n - 1])$ 
13     $\text{CONSTMODMULT}_{1+x^k, m(x)}(H[0..n - 1])$ 

```

for the constant multiplication. In a strict comparison of these CNOT gates, this is worse than the $O(n)$ CNOT gates used by modular schoolbook multiplication, even if we can find a better estimate, but our primary goal is minimizing the number of Toffoli gates without introducing ancillary qubits. In our implementation, even the sum of CNOT and Toffoli gates ends up lower after some degree than the number of Toffoli gates for schoolbook multiplication.

We implemented Algorithm 11 in Java to simulate the execution. Code can be found in [13]. We used the program to automatically count the number of gates and give an estimate of the depth, see Table 6 for the results. Depth count is done by maintaining a set of gates and checking every gate: if they overlap with the previous gate(s) the depth is increased by

Line	H in $\text{MODMULT}_{m(x)}$
1	h
2-4	$x^{-k}h \mod m$
5-7	$\gamma + x^{-k}h \mod m$
8	$(1 + x^k)^{-1}(\gamma + x^{-k}h) \mod m$
9	$(1 + x^k)^{-1}\gamma + \beta + x^{-k}(1 + x^k)^{-1}h \mod m$
11	$x^k(1 + x^k)^{-1}\gamma + x^k\beta + (1 + x^k)^{-1}h \mod m$
12	$\alpha + x^k(1 + x^k)^{-1}\gamma + x^k\beta + (1 + x^k)^{-1}h \mod m$
13	$(1 + x^k)\alpha + x^k\gamma + x^k(1 + x^k)\beta + h \mod m$

Table 5. Step-by-step calculation of Algorithm 6.

Degree	schoolbook TOF gates	Algorithm 11 TOF gates	CNOT gates	Depth
2	4	3	9	9
4	16	9	44	32
8	64	27	200	124
16	256	81	678	365
32	1,024	243	2,238	1,110
64	4,096	729	6,896	3,129
127	16,129	2,185	20,632	8,769
128	16,384	2,187	21,272	9,142
163	26,569	4,387	37,168	17,906
233	54,289	6,323	63,655	29,530
256	65,536	6,561	64,706	26,725
283	80,089	10,273	89,620	41,548
571	326,041	31,171	270,940	121,821
1024	1,048,576	59,049	591,942	234,053

Table 6. CNOT and TOF gate count and depth upper bounds for various instances of Algorithm 11 as well as TOF gate count for schoolbook multiplication. Field polynomials used are the same as in Table 1, with the irreducible polynomial chosen that has the lowest CNOT count. Depth numbers are upper bounds from the java implementation.

1 and if they are not overlapping the gate is added to the set of gates to check against. The set of gates is cleared and replaced with the last gate whenever the depth is increased. The author is aware of methods to improve the depth but leaves this to future work.

When doing classical Karatsuba multiplication, the recursive Karatsuba multiplication is often substituted for schoolbook multiplication starting at a cutoff. For example, if multiplication is at most 7 times as expensive as addition, multiplication of two polynomials of degree at most 2 might be replaced by schoolbook multiplication to get 4 TOF gates instead of 3 TOF and 8 CNOT gates. However, the author is unaware of any realistic estimates of cost difference between CNOT and Toffoli gates where the difference is this small.

6.1 Comparison to other instances of binary finite field multiplication

\mathbb{F}_{2^n} $n =$	Toffoli gates			CNOT gates			qubits		
	Here	[7]	[14]	Here	[7]	[14]	Here	[7]	[14]
4	9	9	16	44	22	3	12	17	12
16	81	81	256	678	376	45	48	113	48
127	2185	2185	16129	20632	13046	126	381	2433	381
256	6561	6561	65536	64706	57008	765	768	7073	768
n	$O(n^{\log_2 3})$	$O(n^{\log_2 3})$	n^2	$O(n^2)$	$O(n^{\log_2 3})$	$O(n)$	$3n$	$O(n^{\log_2 3})$	$3n$

Table 7. Comparison of this work with the works of Kepley and Steinwandt [7] and Maslov et al. [14] in terms of Toffoli and CNOT gates as well as qubit count.

We compare our algorithm to two previous instances of multiplication: a variant by Kepley and Steinwandt [7] that optimizes TOF gate count and a variant by Maslov, Mathew, Cheung and Pradhan [14] that does not use Karatsuba. Other variants exist, such as a Karatsuba variant by Parent, Roetteler and Mosca [6], that are worse in terms of space or Toffoli gate count. Since Kepley and Steinwandt use Clifford and T-gates rather than CNOT and Toffoli,

we translate 7 of their T-gates and 8 Clifford gates to 1 Toffoli gate, and translate any remaining Clifford gates to CNOT. The resulting comparison is in Table 7. We can see that although Algorithm 11 does not compare favorably in every regard, both the number of Toffoli gates and the number of qubits are best compared to the alternatives.

7 Conclusion

Algorithm 11 provides a multiplication algorithm for binary polynomials in finite fields without using ancillary qubits and which has sub-quadratic Toffoli gate count. The CNOT gate count is high and the depth is not optimized, which is left open for future work: multiplication by a constant polynomial in \mathbb{F}_{2^n} can likely be done in approximately linear time, which would bring down the theoretical CNOT gate count to the same order as classical Karatsuba. The saving in Toffoli gate count is the same as for Karatsuba on classical computers: for cryptographic field sizes the savings in Toffoli gates ranges from 80 to over 90 percent. This provides a basis for future work on elliptic curve problems on quantum computers as well as potential other work.

Acknowledgements

The author thanks Tanja Lange for her insights into quantum algorithms and classical finite field operations, Tanja Lange and Gustavo Banegas for their advice and supervision both on this paper and the master thesis this paper originates from, and to Daniel J. Bernstein for his insights into both quantum computing and classical multiplication algorithms. The author also thanks the referees for providing valuable feedback on the paper.

References

1. A. A. Karatsuba and Y. P. Ofman (1962), *Multiplication of many-digit numbers by automatic computers*, Doklady Akademii Nauk, vol. 145, pp. 293–294.
2. R. Maeder (1993), *Storage Allocation for the Karatsuba Integer Multipliation Algorithm*, Design and Implementation of Symbolic Computation Systems, International Symposium, DISCO '93, Gmunden, Austria, September 15-17, 1993, Proceedings, pp. 59–65.
3. E. Thomé (2002), *Karatsuba multiplication with temporary space of size $\leq n$* , <https://members.loria.fr/ETHome/files/kara.pdf>.
4. D. S. Roche (2009), *Space- and Time-Efficient Polynomial Multiplication*, Proceedings of the 2009 international symposium on Symbolic and algebraic computation, pp. 295–302.
5. Y. Cheng (2016), *Space-Efficient Karatsuba Multiplication for Multi-Precision Integers*, arXiv preprint arXiv:1605.06760.
6. A. Parent, M. Roetteler, and M. Mosca (2017), *Improved reversible and quantum circuits for Karatsuba-based integer multiplication*, 12th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2017, June 14-16, 2017, Paris, France, pp. 7:1–7:15.
7. S. Kepley and R. Steinwandt (2015), *Quantum circuits for \mathbb{F}_{2^n} -multiplication with subquadratic gate count*, Quantum Information Processing, vol. 14, no. 7, pp. 2373–2386.
8. B. Amento, M. Rötteler, and R. Steinwandt (2013), *Efficient quantum circuits for binary elliptic curve arithmetic: reducing T-gate complexity*, Quantum Information & Computation, vol. 13, no. 7-8, pp. 631–644.
9. R. Avanzi, H. Cohen, C. Doche, G. Frey, T. Lange, K. Nguyen, and F. Vercauteren, eds. (2005), *Handbook of Elliptic and Hyperelliptic Curve Cryptography*, Chapman and Hall/CRC.
10. FIPS (2013), *PUB 186-4: Federal information processing standards publication. digital signature standard (DSS)*, Information Technology Laboratory, National Institute of Standards and Tech-

- nology (NIST), Gaithersburg, MD.
11. G. Banegas, R. Custódio, and D. Panario (2018), *A new class of irreducible pentanomials for polynomial-based multipliers in binary fields*, Journal of Cryptographic Engineering, pp. 1–15.
 12. G. Seroussi (1998), *Table of low-weight binary irreducible polynomials*, Hewlett-Packard Laboratories.
 13. I. van Hoof (2019), *QMKMBP: Quantum modulo Karatsuba multiplier for binary polynomials*, Github, <https://github.com/ikbenbeter/QMKMBP>.
 14. D. Maslov, J. Mathew, D. Cheung, and D. K. Pradhan (2009), *An $O(m^2)$ -depth quantum algorithm for the elliptic curve discrete logarithm problem over $GF(2^m)^a$* , Quantum Information & Computation, vol. 9, no. 7, pp. 610–621.