



Efficiently Masking Polynomial Inversion at Arbitrary Order

Markus Krausz¹, Georg Land^{1,2(✉)}, Jan Richter-Brockmann¹,
and Tim Güneysu^{1,2}

¹ Horst Görtz Institute for IT Security, Ruhr University Bochum, Bochum, Germany

{markus.krausz,georg.land,tim.guneysu}@rub.de

² DFKI GmbH, Cyber-Physical Systems, Bremen, Germany

Abstract. Physical side-channel analysis poses a huge threat to post-quantum cryptographic schemes implemented on embedded devices. Still, secure implementations are missing for many schemes. In this paper, we present an efficient solution for masked polynomial inversion, a main component of the key generation of multiple post-quantum Key Encapsulation Mechanisms (KEMs). For this, we introduce a polynomial-multiplicative masking scheme with efficient arbitrary order conversions from and to additive masking. Furthermore, we show how to integrate polynomial inversion and multiplication into the masking schemes to reduce costs considerably. We demonstrate the performance of our algorithms for two different post-quantum cryptographic schemes on the Cortex-M4. For NTRU, we measure an overhead of 35% for the first-order masked inversion compared to the unmasked inversion while for BIKE the overhead is as little as 11%. Lastly, we verify the security of our algorithms for the first masking order by measuring and performing a TVLA based side-channel analysis.

Keywords: PQC · Masking · Polynomial inversion · Higher-order masking

1 Introduction

Our digital infrastructure relies and trusts Public-Key Cryptography (PKC) to establish secure communication channels. However, due to Shor's algorithm presented in 1999 [36], currently used schemes like RSA [33] and ECC [29] can be broken by quantum computers in polynomial time. Therefore, in 2017, the National Institute of Standards and Technology (NIST) announced a *Post-Quantum Cryptography Standardization Project* to find and standardize new cryptographic schemes that provide security against attacks mounted on classical and quantum computers. After three rounds, the NIST identified seven finalists and eight alternate candidates which are considered for standardization. Besides security, important metrics like costs, performance, and implementation characteristics on var-

M. Krausz, G. Land and J. Brockmann—These authors contributed equally to this work.

ious platforms are considered in the selection process [2]. Driven by these criteria, the research community has proposed a plethora of highly efficient implementations for software and hardware. However, implementations of Post-Quantum Cryptography (PQC) schemes on embedded devices are faced with the same problems as traditional cryptographic algorithms, which includes physical attacks like Side-Channel Analysis (SCA) and Fault-Injection Analysis (FIA).

So far, most of the side-channel research with respect to the finalists in NIST’s PQC standardization process focuses on schemes based on the Learning with Error (LWE) problem. Bos et al. presented the first higher-order masked implementation for the Cortex-M0+ and the Cortex-M4 for Kyber [8]. Just recently, Heinz et al. published a report on an optimized first-order protected Kyber implementation for the Cortex-M4 including practical measurements [19]. In 2021, Beirendonck et al. presented a first-order protected implementation of Saber for the Cortex-M4 [4]. An optimized implementation that also provides protection against higher-order attacks was afterwards proposed in [26].

Besides these studies that directly target the protection of specific algorithms, others [14, 18] proposed optimizations and implementations which can be applied to both schemes. Coron et al. [14] concentrated their work on the improvements of higher-order masked comparisons by considering different approaches and techniques. As a case study, they applied their optimizations to Kyber and Saber. The work of Fritzmann et al. [18] explored different masked accelerators used as instruction set extensions for a RISC-V processor. They demonstrated their improvements on a hardware software co-design for Kyber and Saber. Eventually, D’Anvers et al. improved the work of Coron et al. [14] and presented an optimized higher-order masked comparison [15].

Summarizing, we can see that the side-channel security countermeasures for the LWE problem based schemes Kyber and Saber have already received some attention. However, masking NTRU-like [20, 21] and code-based [3, 28] systems is still an open research question and has so far only been sparsely investigated. In contrast, several side-channel attacks on these schemes were demonstrated. At CHES 2019, Sim et al. present a generic side-channel attack using conditional moves in implementations of PQC schemes based on Quasi-Cyclic Moderate-Density Parity-Check (QC-MDPC) codes [37]. Recently, a single-trace side-channel attack on the polynomial sampling of NTRU, NTRU Prime, and Dilithium has been proposed in [25]. In the work of Mujdei et al. [30] the authors present a powerful correlation power analysis on polynomial multiplications effecting all lattice-based PQC schemes.

An important operation in almost all NTRU-like and code-based systems is the polynomial inversion. It is required in the key generation of the finalists NTRU-HPS and NTRU-HRSS [10] as well as in the two alternate candidates Streamline NTRU Prime [5] and BIKE [3].

Contribution. To this end, we present the first efficient methodology for masking polynomial inversion by introducing polynomial-multiplicative masking (Sect. 3). As a foundation for our approach, we develop secure arbitrary-

order conversions from polynomial-additive to polynomial-multiplicative masking (Sect. 3.1) and vice versa (Sect. 3.2). We show how to integrate a masked polynomial inversion into this conversion to reduce the number of unmasked inversions to one, independent of the masking order (Sect. 3.3). Additionally, we develop an algorithm to integrate a masked polynomial multiplication into the conversion to save costly unmasked multiplications (Sect. 3.4). Finally, we implement our algorithms for two use cases to demonstrate the performance benefits and we back our security claims for the first masking order by performing practical measurements on a Cortex-M4 microcontroller (Sect. 4).

2 Preliminaries

In this section we introduce important preliminaries that are necessary to adequately describe our approaches of masked arithmetic operations. Besides stating notations used throughout this work, we briefly recap masking. Eventually, we describe practical applications of masked polynomial inversions in the field of PQC.

2.1 Notation

Throughout this work, we denote polynomials by x . The i -th share of a shared polynomial x is denoted by x_i . A uniform random sampling of a polynomial r is denoted by $r \stackrel{\$}{\leftarrow} \mathcal{R}$ where \mathcal{R} is the set of all valid polynomials. The set \mathcal{R}^* denotes all uniform sampled polynomials from \mathcal{R} that are invertible.

2.2 Masking

Masking is a common countermeasure to prevent SCA on embedded devices and is studied in the scientific community for more than twenty years [9]. The foundation of masking is *secret sharing* which splits a sensitive value x into multiple shares x_i with $0 \leq i \leq d$. For a correct sharing holds

$$x = x_0 \circ x_1 \circ \dots \circ x_d \tag{1}$$

where \circ defines the group operator of the applied masking scheme and d defines the security order based on the d -probing model proposed in [22]. As a consequence, a function f processing x needs to be transformed as well such that $f = f_0 \circ f_1 \circ \dots \circ f_d$. When applying \oplus as the group operator in Eq. 1, the secret sharing scheme is called *boolean masking*. The encoding is called *arithmetic masking* when \circ is replaced by an addition or multiplication which we further categorize as *additive masking* or *multiplicative masking*, respectively.

2.3 Polynomial Inversion Applications

Polynomial inversion is a regular used operation in several PQC schemes [3, 20, 21]. Since it is such a critical operation, several works concentrated on

efficient implementations of the polynomial inversion for software and hardware [11, 17, 31, 32]. However, most approaches are based on Fermat's Little Theorem performed by the Itoh-Tsujii Algorithm (ITA) algorithm [23] or on the extGCD proposed by Bernstein and Yang [7]. In the following, we will briefly introduce the finalist NTRU, and the two alternate candidates streamlined NTRU Prime and Bit Flipping Key Encapsulation (BIKE) as examples of PQC schemes requiring polynomial inversions.

NTRU. The finalist NTRU is based on the original work by Hoffstein et al. [20] and on the work by Hülsing et al. [21]. NTRU is defined by three coprime positive integers (n, p, q) , the sample spaces $\mathcal{L}_f, \mathcal{L}_g, \mathcal{L}_r, \mathcal{L}_m$, and an injection $\text{Lift} : \mathcal{L}_m \rightarrow \mathbb{Z}[\mathbf{X}]$. Furthermore, the authors of the NTRU submission recommend two families of parameter sets called NTRU-HPS and NTRU-HRSS [10]. NTRU-HPS uses a *fixed-weight* sampling space and allows several choices of q for each n which are based on [20] while NTRU-HRSS uses an *arbitrary weight* sampling space and fixed q as a function of n as suggested in [21].

The key generation requires to perform two polynomial inversions to generate the public and private key as shown in Algorithm 1. Note, for NTRU-HPS as well as for NTRU-HRSS the parameter p is always fixed to three. However, the two parameters (n, q) are different for the three security levels $\lambda \in \{1, 3, 5\}$ and are defined as (509, 2048), (677, 2048), and (821, 4096), respectively.

Streamlined NTRU Prime. Streamlined NTRU Prime [5] is an alternate candidate in the NIST standardization process. NTRU Prime is also based on the original proposal by Hoffstein et al. [20] and defined by a prime number p , a prime number q , and a positive integer w [6]. One of the main differences to the classic NTRU cryptosystem is that NTRU Prime works over prime fields which avoids various attack vectors as claimed by the authors [5]. The key generation in NTRU Prime (see Algorithm 2) also contains two polynomial inversions. The first inversion inverts the randomly sampled polynomial g drawn from R while the second inversion inverts $3 \cdot f$ where f is a polynomial with coefficients $f_i \in \{-1, 0, 1\}$ with exactly w non-zero coefficients. Note, the first sampled polynomial g is not always invertible in R_3 while the second polynomial f is always invertible in R_q since it is a field.

For the three security levels $\lambda \in \{1, 3, 5\}$ the NTRU Prime parameters (p, q, w) are defined as (653, 4621, 288), (953, 6343, 396), and (1277, 7879, 492), respectively.

Algorithm 1. KeyGen NTRU.**Require:** NTRU parameters n, p, q .**Ensure:** Priv. key (f, f_p, f_q) , pub. key h .

- 1: Generate $f \xleftarrow{\$} \mathcal{L}_f$
- 2: Generate $g \xleftarrow{\$} \mathcal{L}_g$
- 3: Compute $f_p \leftarrow 1/f$ in S_3
- 4: Compute $f_q \leftarrow 1/f$ in S_q
- 5: Compute $g \leftarrow 3 \cdot g \cdot f_g$ in R_q
- 6: Compute $h_q \leftarrow 1/h$ in S_q
- 7: Return (f, f_p, f_q) and h .

Algorithm 2. KeyGen sNTRUp.**Require:** sNTRUp parameter q .**Ensure:** Priv. key (f, g_{inv}) , pub. key h .**repeat** Generate $g \xleftarrow{\$} R$, g small**until** g is invertible in R_3 Generate $f \xleftarrow{\$} Short$ Compute $g_{inv} \leftarrow 1/g$ in R_3 Compute $h \leftarrow g/(3 \cdot f)$ in R_q Return (f, g_{inv}) and h .

BIKE. As well as Streamlined NTRU Prime, BIKE has been selected as an alternate candidate. In contrast to NTRU, BIKE is a code-based scheme relying on QC-MDPC codes [3]. The scheme originally consists of three different algorithms BIKE-1, BIKE-2, and BIKE-3 which, however, were reduced to just one single Key Encapsulation Mechanism (KEM) called BIKE. In BIKE, all polynomials are from the cyclic polynomial ring $\mathcal{R} := \mathbb{F}_2[X]/(X^r - 1)$ where r defines the size of the polynomials. The public key h is generated by sampling two private sparse polynomials (h_0, h_1) with $|h_0| = |h_1| = w/2$, inverting h_0 , and multiplying the results with h_1 . The entire key generation is formally described in Algorithm 3. For the three security levels $\lambda \in \{1, 3, 5\}$, the two parameters (r, w) are defined as $(12323, 141)$, $(24659, 206)$, and $(40973, 274)$, respectively. Since BIKE is suggested to be used with ephemeral keys, an efficient masked implementation of the polynomial inversion for side-channel protected designs is necessary.

In summary, it can be seen in Algorithm 1, Algorithm 2, and Algorithm 3 that the polynomial inversion is a major operation in the key generation of all three algorithms. Our measurements in Sect. 4.1 confirm that the polynomial inversion dominates the costs in terms of cycle counts. Hence, to construct protected designs against SCA, it is essential to find efficient algorithms for masked implementations. However, not only the inversion itself should be implemented efficiently but also preceding and subsequent operations must be masked without any expensive conversions between different masking techniques. Before we present our approach of an efficient higher-order masked polynomial inversion, we briefly discuss different cases of invertibility of random polynomials.

Invertibility of Random Polynomials. Among these three schemes, three different cases of invertibility occur. Since the target polynomials are sampled randomly but based on certain rules, we identify the following cases.

Algorithm 3. Key Generation of BIKE.**Require:** BIKE parameters n, w, ℓ .**Ensure:** Private key (h_0, h_1, σ) and public key h .Generate $(h_0, h_1) \xleftarrow{\$} \mathcal{R}^2$ both of odd weight $|h_0| = |h_1| = w/2$.Generate $\sigma \xleftarrow{\$} \{0, 1\}^\ell$ uniformly at random.Compute $h \leftarrow h_1 h_0^{-1}$.Return (h_0, h_1, σ) and h .

1. All sampled polynomials (except the polynomial representing 0) are invertible. This case is trivial and no further exceptions need to be covered which is the case for NTRU.
2. Not all polynomials from the used ring are invertible but following some certain rules always allows to sample an invertible polynomial. For example, this is the case for BIKE where the polynomials requires to have an odd Hamming weight. Hence, applying the defined sampling procedure always generates an invertible polynomials such that the inversion cannot fail.
3. Not all polynomials from the underlying ring are invertible and they are not easily distinguishable. For example, this is the case for Streamlined NTRU Prime where the sampling procedure just sample uniformly random polynomials without applying dedicated rules. In case the sampled polynomial is not invertible, the inversion fails in the last step and a new polynomial needs to be sampled.

3 Masking Polynomial Inversion

Masking boolean operations in PQC schemes can efficiently be implemented with a boolean sharing, while arithmetic operations such as the addition and subtraction of polynomials or the multiplication with public values are implemented with additive sharing as the masked implementation for Kyber [8] demonstrates. An alternative sharing, that had already been proposed for AES in the year 2001 [1], is multiplicative sharing. The problem with multiplicative sharing that hinders its application, is that if one share is zero, the attacker already knows that the masked value is zero.

For polynomial inversion, that is used in multiple PQC schemes as shown in Sect. 2.3, we need a masking approach for which inversion is a linear operation. Given uniformly random polynomials $m_i \in \mathcal{R}$ such that $m = \prod_{i=0}^d m_i$, a valid polynomial-multiplicative sharing can be realized by

$$m^{-1} = \prod_{i=0}^d m_i^{-1}, \quad (2)$$

i.e., the inversion is applied to each share independently. As the zero polynomial is not invertible, it will not be given as an input to a masked inversion. With

$d + 1$ unmasked polynomial inversions, that is already an expensive operation on its own, this approach is very costly and asks for alternative solutions.

Obviously, multiplication of *two secret* polynomials is very efficient in the multiplicative domain as it requires only $d + 1$ unmasked multiplications compared to the additive domain where the number of unmasked multiplications is quadratic to the masking order in current solutions [35]. The cost to convert polynomials from and to the multiplicative domain determines, however, whether this approach is viable (cf. Section 3.4).

In the following, we present algorithms that efficiently transform additive shares of polynomials in a ring to multiplicative shares and vice versa. With the motivation to perform a more efficient polynomial inversion than shown in Eq. 2, we demonstrate how to integrate the inversion into the transformation, and how to perform a multiplication and back transformation in one joint operation.

3.1 Conversion from Additive to Multiplicative Sharing

Let a be a polynomial and a_i shares with $a = \sum_{i=0}^d a_i$, where all a_i are uniform random in the respective polynomial ring. To transform this sharing to a polynomial-multiplicative sharing in the same ring, we adapt the well-known technique of first appending a share in the new masking domain, enlarging the sharing in two domains (additive and multiplicative), and then to combine two old shares to remove one share.

We now introduce our algorithm by presenting an example for first-order masking. Given a polynomial a split into two additive shares a_0 and a_1 , we start by sampling one invertible polynomial r and multiply each additive share with this polynomial, yielding ra_0 and ra_1 . We set the inverted polynomial r^{-1} as a new multiplicative share, expanding the number of shares from two to three. To reduce our number of shares, we add corresponding two additive shares: $ra_0 + ra_1 = r(a_0 + a_1)$. By treating the sum as a multiplicative share, we are left with two correct multiplicative shares for a , since $r^{-1}r(a_0 + a_1) = a$.

The full algorithm for arbitrary orders can be summarized with the following steps:

1. Sample a uniform random and invertible polynomial r , observing that $a = r^{-1}ra$.
2. Compute $a'_i = ra_i$, we now have $d + 2$ shares, $d + 1$ additive shares and one multiplicative share.
3. To return to $d + 1$ shares, we combine two additive shares.
4. Repeat from start until there is only one additive share left, which now can be viewed as a multiplicative share.

The algorithm is shown in detail in Algorithm 4. Note that for this conversion, d polynomial inversions and $(d + 1)(d + 2)/2 - 1$ polynomial multiplications, as well as $d - 1$ polynomial additions are needed.

Algorithm 4. Additive to Polynomial-Multiplicative Masking Conversion (A2M)

```

Require:  $a = \sum_{i=0}^d a_i$ 
Ensure:  $m = \prod_{i=0}^d m_i = a$ 
function A2M( $a_0, \dots, a_d$ )
  for  $i := d$  downto 1 do
     $r \xleftarrow{\$} \mathcal{R}^*$  ▷ sample from the set of invertible polynomials
    for  $j := 0$  to  $i$  do
       $a_j := ra_j$ 
    end for
    end for
     $m_i := r^{-1}$  ▷ now we have  $d + 2$  shares with  $a = \left(\sum_{j=0}^i a_j\right) \prod_{j=i}^d m_j$ 
     $a_{i-1} := a_{i-1} + a_i$  ▷ combining two additive shares
  end for
   $m_0 := a_0$ 
end function

```

3.2 Conversion from Multiplicative to Additive Sharing

For subsequent operations in the additive domain, a transformation from the multiplicative to the additive domain is necessary. Given a masked polynomial m split into two multiplicative shares m_0 and m_1 for our M2A conversion, we start by sampling one random polynomial r . The first step is to compute $m_0 + r$ before we multiply it with m_1 to get $(m_0 + r)m_1 = m_0m_1 + rm_1$. Together with the product $-rm_1$ we have two additive shares that yield $m_0m_1 + rm_1 - rm_1 = m_0m_1 = m$.

This method can be generalized to arbitrary masking orders by reapplying the core idea of adding a random polynomial before the multiplication with one of the multiplicative shares m_i . Our strategy is to compute $m = \prod_{i=0}^d m_i$ step by step in the first share, while protecting this sum with d random summands. Thus, iterating from $i = 1$ to d , we sample a uniform random additive sharing of $i + 1$ polynomials such that $\sum_{j=0}^i r_{ij} = 0$. We add these random polynomials to the first $i + 1$ shares before we multiply the shares with m_i . After d iterations, we get $a_0 = m + \sum_{i=1}^d (r_{i0} \prod_{j=i}^d m_j)$ as the first additive share for m together with d additive shares $a_k = \sum_{i=k}^d (r_{ik} \prod_{j=i}^d m_j)$ that cancel out the summands in a_0 except m .

The algorithm can efficiently be implemented in situ as shown in Algorithm 5 and utilizes $d(d + 1)/2 + d$ polynomial multiplications, $d(d + 1) + d$ additions, $d(d + 1)/2$ fresh random polynomials and no costly inversion.

3.3 Reducing the Number of Inversions

The main application of the polynomial-multiplicative masking is polynomial inversion. Naively, we would perform a polynomial inversion on each polynomial-multiplicative share individually to obtain a sharing of the inverted polynomial (cf. Equation 2). Together with the d inversions required for the A2M conversion,

Algorithm 5. Polynomial-Multiplicative to Additive Masking Conversion (M2A)

Require: $m = \prod_{i=0}^d m_i$
Ensure: $a = \sum_{i=0}^d a_i = m$
function M2A(m_0, \dots, m_d)
 for $i := 1$ **to** d **do**
 $r_i := 0$
 for $j := 0$ **to** $i - 1$ **do**
 $r \xleftarrow{\$} \mathcal{R}$
 $r_i := r_i + r$
 $m_j := m_j + r$ ▷ refreshing
 $m_j := m_j m_i$ ▷ combining two multiplicative shares
 end for
 $m_i := -r_i m_i$
 end for
 for $i := 0$ **to** d **do**
 $a_i := m_i$
 end for
end function

this would lead to $2d + 1$ unmasked inversions for one masked inversion, given a polynomial shared in the additive domain.

However, we can adapt Algorithm 4 such that only one polynomial inversion is necessary, independent of the masking degree. This is shown in Algorithm 6. The idea is to not set the new multiplicative shares to the inverse, which we would invert again later, but to the original sample. Instead we only invert m_0 at the end to get an A2M conversion with implicit inversion. With this method we can drastically reduce the number of polynomial inversions that are the most expensive operations compared to polynomial multiplications and additions as we show in Sect. 4. We thus save two inversions for first order, four inversions for second and already six inversions for third order masking, compared to the naive approach.

3.4 Reducing the Number of Multiplications

Although a masked polynomial multiplication is cheaper in the multiplicative domain ($d + 1$ unmasked multiplications) compared to the additive domain where it is quadratic [35], the additional costs of the A2M and M2A conversions render this approach obsolete for polynomials that are not given in the multiplicative domain anyway. In particular the A2M conversion without inversion is too expensive with its d unmasked inversions.

We can, however, save unmasked multiplications when one factor is already in the multiplicative domain due to a prior inversion. Given a polynomial $a = \sum_{i=0}^d a_i$ in the additive domain and a polynomial $b = \prod_{i=0}^d b_i$ in the multiplicative domain, we observe that the masked product $c = \sum_{i=0}^d c_i = ab$

Algorithm 6. Additive to Polynomial-Multiplicative Masking Conversion with Implicit Polynomial Inversion ($A2M_{INV}$)

Require: $a = \sum_{i=0}^d a_i$
Ensure: $m = \prod_{i=0}^d m_i = a^{-1}$
function $A2M_{INV}(a_0, \dots, a_d)$
 for $i := d$ **downto** 1 **do**
 $r \xleftarrow{\$} \mathcal{R}^*$ ▷ sample from the set of invertible polynomials
 for $j := 0$ **to** i **do**
 $a_j := ra_j$
 end for
 $m_i := r$ ▷ note that we do not set this to the inverse
 $a_{i-1} := a_{i-1} + a_i$
 end for
 $m_0 := a_0^{-1}$ ▷ the only inverse we need to compute
end function

can be computed with $c = ab = \sum_{i=0}^d a_i \prod_{j=0}^d b_j = \sum_{i=0}^d (a_i \prod_{j=0}^d b_j)$, where $c_i = a_i \prod_{j=0}^d b_j$ represents an additive share of the product c . The straightforward computation would leak the polynomial b , but by adding fresh random polynomials between the unmasked multiplications similar as in our $M2A$ conversion, we can get a secure conversion from multiplicative domain to additive domain including a multiplication with an additive shared polynomial as shown in Algorithm 7.

The costs for this masked conversion with implicit multiplication are $(d + 1)^2$ unmasked multiplications, $(d + 1)2d$ additions and $(d + 1)d$ fresh random polynomials. Compared to the naive approach of first converting a from the multiplicative to the additive domain and then performing the multiplication, we save about the amount of unmasked multiplications and additions required for the $M2A$ conversion.

For the case where we want to securely invert a polynomial and multiply the result with another polynomial, which is often the case as we see in Sect. 2.3, we apply our $A2M_{INV}$ first, where the costs are dominated by the single unmasked inversion, resulting in an inverted polynomial in the multiplicative domain. As a second step, we apply our $M2A_{MUL}$, to transform the inverted polynomial back into the additive domain while simultaneously multiplying it with another additive shared polynomial, at the cost of a multiplication in the additive domain, so the back transformation is basically free. In Sect. 4 we present performance results by exemplary applying our approaches to NTRU and BIKE.

4 Implementation and Evaluation

To evaluate the performance and security of our algorithms, we implemented them for NTRU and BIKE on the STM32F4 discovery board, which is equipped with a 32-bit Cortex-M4 microcontroller, 192-KB SRAM and 1-MB flash memory and can be clocked up to 168 MHz.

Algorithm 7. Polynomial-Multiplicative to Additive Masking Conversion with Implicit Polynomial Multiplication ($M2A_{MUL}$)

Require: $a = \sum_{i=0}^d a_i, b = \prod_{i=0}^d b_i$
Ensure: $c = \sum_{i=0}^d c_i = ab$
function $M2A_{MUL}(a_0, \dots, a_d, b_0, \dots, b_d)$
 for $j := 0$ **to** d **do**
 $c_j := a_j b_0$ ▷ implicit multiplication
 end for
 for $i := 1$ **to** d **do**
 $r_d := 0$
 for $j := 0$ **to** $d - 1$ **do**
 $r \xleftarrow{\$} \mathcal{R}$
 $r_d := r_d + r$
 $c_j := c_j + r$ ▷ refreshing
 $c_j := c_j b_i$ ▷ combining two multiplicative shares
 end for
 $c_d := c_d - r_d$
 $c_d := c_d b_i$
 end for
end function

We based our implementation on the respective ring operations of the state-of-the-art Cortex-M4 implementations of the schemes. For BIKE this is the work by Chen et al. [12], for NTRU this is the work by Chung et al. [13] with an improved inversion by Li et al. [27].

4.1 Implementation Results

As it is common [24], we measured cycle counts at 24 MHz to not have memory wait states. We compiled our code with the `arm-none-eabi-gcc-10.3.1` compiler with optimization-level `O3`. The stated cycle counts are averages of 100 runs.

We did not implement and measure the plain A2M conversion, because it is not interesting for our use cases with its high costs.

NTRU. We first measured the cycle counts for unmasked ring operations to have a baseline to compare our masked versions with. For NTRU in the parameter set `ntruhs2048677`, polynomials in the ring S_3 have 677 coefficients $\in \{0, 1, 2\}$. Unprotected polynomial inversion costs 1273864 clock cycles here, about six times the cycles for an unprotected polynomial multiplication that takes 201383 cycles. An unprotected addition is done in only 18340 cycles and is thus insignificant compared to inversions and multiplications.

The costs for the masked $A2M_{INV}$ in the first masking order are mainly determined by the unmasked inversion and two unmasked multiplications. The overhead compared to an unmasked inversion is therefore mostly the cost of two

multiplications, resulting in about 35% overhead, which is an excellent result compared to other masked operations. This calculation excludes the cost of an M2A conversion, but as we argued in Sect. 3.4, this comes for free by using the $M2A_{MUL}$. Since the number of unmasked inversions required for one $A2M_{INV}$ is only one, independent of the masking order, the cycle counts of the $A2M_{INV}$ increase only slowly with the masking order. For the sixth order, which operates on seven shares, the cycle counts are less than six fold the ones of the unmasked as shown in Table 1.

For the $M2A_{MUL}$ we measured 885773 cycles in the first order, less than twice the cost of one M2A that costs 486165. This proportion stays with increasing masking order while the number of unmasked multiplications and additions grows quadratically for both algorithms.

Table 1. Cycle counts for our proposed masked $A2M_{INV}$, $M2A_{MUL}$, and M2A conversion for ntruhs2048677 on the Cortex-M4. Unprotected addition requires 18340 clock cycles, unprotected multiplication requires 201383 clock cycles and unprotected inversion 1273864 clock cycles.

Ord. d	A2M inversion			M2A mul.			M2A conversion		
	Cycles	MUL	INV	Cycles	MUL	ADD	Cycles	MUL	ADD
1	1723778	2	1	885773	4	4	486165	2	3
2	2372502	5	1	2090841	9	12	1230767	5	8
3	3211410	9	1	3802004	16	24	2238833	9	15
4	4260732	14	1	6057128	25	40	3503189	14	24
5	5524861	20	1	8848501	36	60	5049140	20	35
6	6991050	27	1	12097869	49	84	6859272	27	48

BIKE. For BIKE in the parameter set `bike11`, polynomials have 12323 coefficients $\in \{0, 1\}$. As 32 coefficients are stored in one register and the addition of coefficients equates to a xor operation, the unmasked addition of polynomials is very cheap with 3534 clock cycles. Due to the higher polynomial degree, however, multiplications and inversions take longer, compared to the operations in NTRU. For one unmasked multiplication, we measured about one million cycles, and for one unmasked inversion 19182916 cycles.

With the increased gap between multiplication and inversion, compared to NTRU, the overhead of the $A2M_{INV}$ reduces. With 21317392 cycles for the first order $A2M_{INV}$, the overhead is as little as 11% compared to an unmasked inversion. Also the cost of $M2A_{MUL}$ and M2A become less significant compared a $A2M_{INV}$ in the lower masking orders, due to the order of magnitude difference in cycle counts between unmasked inversion and unmasked multiplication. In the first masking order we measure 4240017 cycles for one $M2A_{MUL}$ and 2131405 for one M2A as shown in Table 2. The gap between $A2M_{INV}$ and M2A or $M2A_{MUL}$

decreases in relative terms with increasing masking order due to the quadratic cost in unmasked multiplications.

Table 2. Cycle counts for our proposed masked $A2M_{INV}$, $M2A_{MUL}$, and $M2A$ conversion for `bikel1` on the Cortex-M4. Unprotected addition requires 3534 clock cycles, unprotected multiplication requires 1052253 clock cycles and unprotected inversion 19182916 clock cycles.

Ord. d	A2M inversion			M2A mul.			M2A conversion		
	Cycles	MUL	INV	Cycles	MUL	ADD	Cycles	MUL	ADD
1	21317392	2	1	4240017	4	4	2131405	2	3
2	24487146	5	1	9584999	9	12	5342630	5	8
3	28736397	9	1	17068753	16	24	9622491	9	15
4	34007250	14	1	26740596	25	40	14994627	14	24
5	40275530	20	1	38507790	36	60	21419851	20	35
6	47744390	27	1	52493255	49	84	28945019	27	48

4.2 Side-Channel Evaluation

To evaluate the security against power side-channel attacks, we performed measurements on the same STM32F4 discovery board with the Cortex-M4 microcontroller. The power consumption is indirectly measured via a $1\ \Omega$ shunt resistor placed in the supply path of the microcontroller (the board provides dedicated pads for such applications) and the signal is amplified by a ZFL-1000LN+ Low Noise Amplifier (LNA). We use an 8 bit oscilloscope from PicoScope sampling with 625 MS/s to acquire the power traces. During the measurements, the microcontroller operates with a 24 MHz clock, which results in roughly 26 sample points per clock cycle, and is powered by an external power supply to ensure a clean and stable supply voltage.

For the security evaluation, we use a common fixed vs. random univariate Test Vector Leakage Assessment (TVLA) evaluation procedure as detailed described in [34]. Commonly, the measured power traces of the fixed and random inputs are used for a Welsh t -test where the t -value is compared to a ± 4.5 threshold corresponding to a $\alpha = 0.0001$ confidence level. In case the threshold is exceeded, the implementation is assumed to leak sensitive information since the power consumption of the fixed and the random inputs can be distinguished. However, in 2017 Ding et al. demonstrated that the threshold of ± 4.5 needs to be adapted for measurements with many sample points to avoid false positives in the evaluation [16]. Since we measure operations that require up to $1.7e6$ clock cycles (which are approximately $26 \cdot 1.7e6 = 44.2e6$ sample points with our setup), we applied their approach and adapted the corresponding threshold that still results in a confidence level of α .

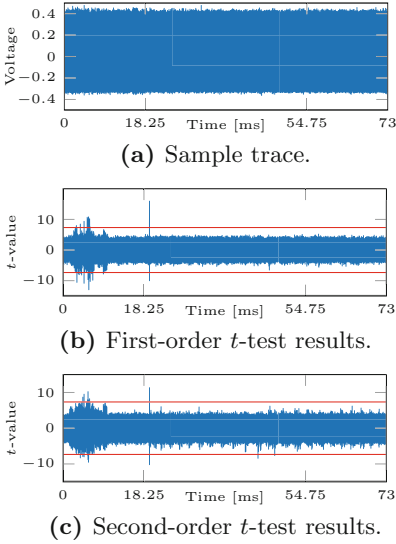


Fig. 1. Measurement results of $A2M_{INV}$ with no randomness (2000 traces).

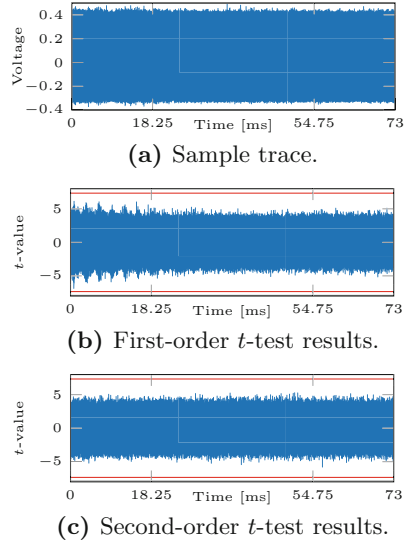



Fig. 2. Measurement results of $A2M_{INV}$ with randomness (100000 traces).

In the following, we present the measurement results for the first-order masked inversion $A2M_{INV}$ and the multiplicative to additive conversion $M2A$. We limit our evaluation to these two algorithms as they exemplarily demonstrate the ideas of our proposals. Both, the $A2M$ conversion and the $M2A_{MUL}$, are similar to the other two algorithms such that we only performed the time-consuming measurements for them.

Masked Inversion. Figure 1 shows the measurement results for the masked inversion presented in Algorithm 6 with disabled randomness to demonstrate the correct functionality of our measurement setup. As expected, the t -test reveals first- and second-order univariate leakage. Figure 2 presents the measurement results for the protected inversion with randomness enabled. We acquired 100000 power traces and could not detect any first-order univariate leakage. Interestingly, the second-order t -test also does not reveal any leakage which is may due to the univariate analysis technique applied in our evaluation. We expect that second-order leakage would be visible once an attacker utilizes multivariate analysis techniques, i.e., combines samples from multiple points in time. Another reason for this phenomena could be the applied masking technique. When we look at a single coefficient of a polynomial with multiplicative sharing, it can not be recreated by $d+1$ respective coefficients of the polynomial shares, but depends on other coefficients too. For the first masking order we combine one random coefficient of one polynomial with all random coefficients of another polynomial



(a) Sample trace.




(b) First-order t -test results.(c) Second-order t -test results.

Fig. 3. Measurement results of the M2A conversion with no randomness (2000 traces).



(a) Sample trace.



(b) First-order t -test results.(c) Second-order t -test results.

Fig. 4. Measurement results of the M2A conversion with randomness (100000 traces).

which can be seen as some kind of higher-order masking. However, this artifact is out of scope of this work and we leave the investigation for future work.

Multiplicative to Additive Conversion. Besides the masked polynomial inversion, we additionally evaluate the multiplicative to additive conversion M2A from Algorithm 5. Again, we first measured the operation with disabled randomness (masks and fresh randomness are constant) which is visualized in Fig. 3. After 2000 traces, the t -test results for the first- and second-order clearly indicate leakage. However, in the next experiment we enable all randomness and perform 100000 measurements. The t -test does not reveal any leakage which is shown in Fig. 4. Again, no second-order leakage is visible due to the same argumentation as above.

5 Conclusion

In this work, we demonstrate that polynomial-multiplicative sharing is a viable solution to mask arithmetic operations of multiple PQC schemes. To this end, we propose an efficient higher-order masked polynomial inversion with implicit additive to multiplicative conversion, conversion algorithms used to switch between different sharings, and a novel masked multiplication that accepts an additive

shared operand and a multiplicative shared operand. Applying our masked polynomial inversion to NTRU, the first-order masked design requires an overhead of only 35%, while the overhead for BIKE is only 11%.

However, there are still masking solutions missing for other operations to have all the pieces necessary for a masked implementation of NTRU or BIKE, which is an interesting target for future work. Another open question is the additional security that polynomial-multiplicative masking provides, when looking at the coefficient level. As already mentioned in Sect. 4.2, traditional masking schemes split one value into $d + 1$ values. But in polynomial multiplication, all coefficients are combined with each other and make one coefficient of the masked polynomial dependent of more than $d + 1$ values.

Acknowledgments. This work was supported by the German Research Foundation under Germany’s Excellence Strategy – EXC 2092 CASA – 390781972, through the H2020 project PROMETHEUS (grant agreement ID 780701), and by the Federal Ministry of Education and Research of Germany through the QuantumRISC (16KIS1038) and PQC4Med (16KIS1044) projects.

References

1. Akkar, M.-L., Giraud, C.: An implementation of DES and AES, secure against some attacks. In: Koç, Ç.K., Naccache, D., Paar, C. (eds.) CHES 2001. LNCS, vol. 2162, pp. 309–318. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44709-1_26
2. Alagic, G., et al.: Status report on the first round of the NIST post-quantum cryptography standardization process. US Department of Commerce, National Institute of Standards and Technology (2019). https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=927303
3. Aragon, N., et al.: BIKE: bit flipping key encapsulation (2021). https://bikesuite.org/files/v4.2/BIKE_Spec.2021.07.26.1.pdf
4. Van Beirendonck, M., D’anvers, J.-P., Karmakar, A., Balasch, J., Verbauwhede, I.: A Side-channel Resistant Implementation of SABER. *ACM J. Emerg. Technol. Comput. Syst. (JETC)* **17**(2), 1–26 (2021)
5. Bernstein, D.J., Chuengsatiansup, C., Lange, T., van Vredendaal, C.: NTRU prime: reducing attack surface at low cost. In: Adams, C., Camenisch, J. (eds.) SAC 2017. LNCS, vol. 10719, pp. 235–260. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-72565-9_12
6. Bernstein, D.J., Chuengsatiansup, C., Lange, T., van Vredendaal, C.: Ntru prime: round 3. Submission to the NIST PQC standardization process (2020). <https://ntruprime.cr.yp.to>
7. Bernstein, D.J., Yang, B.-Y.: Fast constant-time GCD computation and modular inversion. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* **2019**(3), 340–398 (2019)
8. Bos, J.W., Gourjon, M., Renes, J., Schneider, T., van Vredendaal, C.: Masking kyber: first- and higher-order implementations. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* **2021**(4), 173–214 (2021)
9. Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards sound approaches to counteract power-analysis attacks. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 398–412. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48405-1_26

10. Chen, C., et al.: NTRU - algorithm specifications and supporting documentation. Brown University and Onboard security company, Wilmington USA (2019)
11. Chen, M.-S., Chou, T.: Classic McEliece on the arm cortex-M4. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* **2021**(3), 125–148 (2021)
12. Chen, M.S., Güneysu, T., Krausz, M., Thoma, J.P.: Carry-less to BIKE faster. In: Ateniese, G., Venturi, D. (eds.) *Applied Cryptography and Network Security - 20th International Conference, ACNS 2022, Rome, Italy, 20–23 June 2022, Proceedings*, vol. 13269 of *Lecture Notes in Computer Science*, pp. 833–852. Springer, Heidelebrg (2022). https://doi.org/10.1007/978-3-031-09234-3_41
13. Chung, C.M.M., Hwang, V., Kannwischer, M.J., Seiler, G., Shih, C.J., Yang, B.Y.: NTT multiplication for NTT-unfriendly rings new speed records for saber and NTRU on cortex-M4 and AVX2. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* **2021**(2), 159–188 (2021)
14. Coron, J.S., Gérard, F., Montoya, S., Zeitoun, R.: High-order polynomial comparison and masking lattice-based encryption. *Cryptology ePrint Archive* (2021)
15. D’Anvers, J.P., Van Beirendonck, M., Verbauwhede, I.: Revisiting higher-order masked comparison for lattice-based cryptography: algorithms and bit-sliced implementations. *IACR Cryptol. ePrint Arch.*, p. 110 (2022)
16. Ding, A.A., Zhang, L., Durvaux, F., Standaert, F.-X., Fei, Y.: Towards sound and optimal leakage detection procedure. In: Eisenbarth, T., Teglia, Y. (eds.) *CARDIS 2017. LNCS*, vol. 10728, pp. 105–122. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75208-2_7
17. Drucker, N., Gueron, S., Kostic, D.: Fast polynomial inversion for post quantum QC-MDPC cryptography. In: Dolev, S., Kolesnikov, V., Lodha, S., Weiss, G. (eds.) *CSCML 2020. LNCS*, vol. 12161, pp. 110–127. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-49785-9_8
18. Fritzmann, T., et al.: Masked accelerators and instruction set extensions for post-quantum cryptography. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* **2022**(1), 414–460 (2021)
19. Heinz, D., Kannwischer, M.J., Land, G., Pöppelmann, T., Schwabe, P., Sprenkels, D.: First-order masked kyber on ARM cortex-M4. *Cryptology ePrint Archive, Report 2022/058* (2022). <https://ia.cr/2022/058>
20. Hoffstein, J., Pipher, J., Silverman, J.H.: NTRU: a ring-based public key cryptosystem. In: Buhler, J.P. (ed.) *ANTS 1998. LNCS*, vol. 1423, pp. 267–288. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0054868>
21. Hülsing, A., Rijneveld, J., Schanck, J., Schwabe, P.: High-speed key encapsulation from NTRU. In: Fischer, W., Homma, N. (eds.) *CHES 2017. LNCS*, vol. 10529, pp. 232–252. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66787-4_12
22. Ishai, Y., Sahai, A., Wagner, D.: Private circuits: securing hardware against probing attacks. In: Boneh, D. (ed.) *CRYPTO 2003. LNCS*, vol. 2729, pp. 463–481. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45146-4_27
23. Itoh, T., Tsujii, S.: A fast algorithm for computing multiplicative inverses in $GF(2^m)$ using normal bases. *Inf. Comput.* **78**(3), 171–177 (1988)
24. Kannwischer, M.J., Rijneveld, J., Schwabe, P., Stoffelen, K.: PQM4: post-quantum crypto library for the ARM cortex-M4. <https://github.com/mupq/pqm4>
25. Karabulut, E., Alkim, E., Aysu, A.: Single-trace side-channel attacks on ω -small polynomial sampling: with applications to NTRU, NTRU prime, and crystals-dilithium. In: *HOST*, pp. 35–45. IEEE (2021)
26. Kundu, S., D’Anvers, J.P., Van Beirendonck, M., Karmakar, A., Verbauwhede, I.: Higher-order masked Saber. *IACR Cryptol. ePrint Arch.*, 389 (2022)

27. Li, C.L.: Implementation of polynomial modular inversion in lattice based cryptography on ARM (2021)
28. Melchor, C.A., et al.: Hamming Quasi-Cyclic (HQC) - Third round version (2021)
29. Miller, V.S.: Use of elliptic curves in cryptography. In: Williams, H.C. (ed.) CRYPTO 1985. LNCS, vol. 218, pp. 417–426. Springer, Heidelberg (1986). https://doi.org/10.1007/3-540-39799-X_31
30. Mujdei, C., et al.: Side-channel analysis of lattice-based post-quantum cryptography: exploiting polynomial multiplication. IACR Cryptol. ePrint Arch., 474 (2022)
31. Richter-Brockmann, J., Chen, M.-S., Ghosh, S., Güneysu, T.: Racing BIKE: improved polynomial multiplication and inversion in hardware. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2022**(1), 557–588 (2022)
32. Richter-Brockmann, J., Mono, J., Güneysu, T.: Folding BIKE: scalable hardware implementation for reconfigurable devices. IEEE Trans. Comput. **71**(5), 1204–1215 (2022)
33. Rivest, R.L., Shamir, A., Adleman, L.: A method for obtaining digital signatures and public-key cryptosystems. Commun. ACM **21**(2), 120–126 (1978)
34. Schneider, T., Moradi, A.: Leakage assessment methodology - a clear roadmap for side-channel evaluations. In: Güneysu, T., Handschuh, H. (eds.) CHES 2015. LNCS, vol. 9293, pp. 495–513. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48324-4_25
35. Schneider, T., Paglialonga, C., Oder, T., Güneysu, T.: Efficiently masking binomial sampling at arbitrary orders for lattice-based crypto. In: Lin, D., Sako, K. (eds.) PKC 2019. LNCS, vol. 11443, pp. 534–564. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17259-6_18
36. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Rev. **41**(2), 303–332 (1999)
37. Sim, B.Y., Kwon, J., Choi, K.Y., Cho, J., Park, A., Han, D.G.: Novel side-channel attacks on quasi-cyclic code-based cryptography. IACR Trans. Cryptogr. Hardw. Embed. Syst., 180–212 (2019)