

Defeating Low-Cost Countermeasures against Side-Channel Attacks in Lattice-based Encryption

A Case Study on Crystals-Kyber

Prasanna Ravi¹, Thales Paiva², Dirmanto Jap¹, Jan-Pieter D’Anvers³ and Shivam Bhasin¹

¹ Temasek Laboratories, Nanyang Technological University, Singapore

² Fundep and CASNAV, Brazil

³ imec-COSIC KU Leuven, Kasteelpark Arenberg 10 - bus 2452, 3001 Leuven, Belgium

prasanna.ravi@ntu.edu.sg thalespaiva@gmail.com djap@ntu.edu.sg

janpieter.danvers@esat.kuleuven.be sbhasin@ntu.edu.sg

Abstract. In an effort to circumvent the high cost of standard countermeasures against side-channel attacks in post-quantum cryptography, some works have developed low-cost detection-based countermeasures. These countermeasures try to detect maliciously generated input ciphertexts and react to them by discarding the ciphertext or secret key. In this work, we take a look at two previously proposed low-cost countermeasures: the ciphertext sanity check and the decapsulation failure check, and demonstrate successful attacks on these schemes. We show that the first countermeasure can be broken with little to no overhead, while the second countermeasure requires a more elaborate attack strategy that relies on valid chosen ciphertexts. Thus, in this work, we propose the first chosen-ciphertext based side-channel attack that only relies on valid ciphertexts for key recovery. As part of this attack, a third contribution of our paper is an improved solver that retrieves the secret key from linear inequalities constructed using side-channel leakage from the decryption procedure. Our solver is an improvement over the state-of-the-art Belief Propagation solvers by Pessl and Prokop, and later Delvaux. Our method is simpler, easier to understand and has lower computational complexity, while needing less than half the inequalities compared to previous methods.

Keywords: Lattice-based cryptography · Side-Channel Attack · Kyber · Key Encapsulation Mechanism · Chosen Ciphertext Attack

1 Introduction

In July 2022, a new Key Encapsulation Mechanism (KEM) and 3 digital signature schemes were selected as the first standardized PQC algorithms [AAC⁺22], a result of an extensive NIST Post-Quantum standardization process. Kyber KEM [ABD⁺20], based on the well-known Module Learning With Error (MLWE) problem, was the only algorithm standardized for KEMs, and we will soon witness wide-scale adoption of Kyber on a variety of computing devices, including resource-constrained platforms such as embedded microcontrollers [KRSS19, AHKS22]. This naturally makes them susceptible to side-channel attacks, which was also an important consideration during the NIST PQC standardization process [RCDB22]. Thus, several works were reported on SCA of Kyber KEM and proposals for concrete countermeasures [BGR⁺21, RRD⁺23, RRCB20].

Looking at these attacks, the decapsulation procedure of Kyber KEM typically serves as the main target, and they can all be categorized into two broad categories: (1) Known

Ciphertext (KC) based SCA and (2) Chosen Ciphertext based SCA. In KC-based SCA, the attacker only requires the knowledge of inputs to the decapsulation procedure [PPM17, PP19, MWK⁺22]. However, for CC-based SCA, the attacker must control the input to the decapsulation procedure [RRCB20, XPRO20, BDH⁺21]. CC-based attacks form the largest category of attacks and are considered more powerful than KC attacks. The reason is that CC-based attacks can exploit leakage from the entire decapsulation procedure for key recovery, while KC-based attacks have a much narrower attack surface for key recovery. This means that protection against CC-based attacks is costlier than KC-based attacks and requires combining masking and shuffling countermeasures [BGR⁺21, RPBC20]. Thus, CC-based side-channel attacks form the main focus of our work.

CC-based side-channel attacks utilize maliciously crafted ciphertexts for key recovery. Standard countermeasures against these attacks include masking and shuffling, which come at a certain cost. In a different approach, some works have focused on low-cost detection-based countermeasures [XPRO20, RCDB22]. The idea of detection-based countermeasures is to look for maliciously formed ciphertexts. Upon detection of a malicious ciphertext, the ciphertext can be discarded, or the secret key can be refreshed, preventing further exposure towards the attacker. While these countermeasures look attractive for a designer owing to their low cost, we are unaware of a dedicated study of the concrete protection they offer against CC-based side-channel attacks. In this work, we attempt to bridge this gap, we perform a concrete analysis of two such detection countermeasures against CC-based attacks, the (1) Ciphertext Sanity Check and the (2) Decapsulation Failure Check, and demonstrate new types of CC-based attacks are capable of breaking both these countermeasures.

Our Contribution

In this work, we present the following new contributions:

1. We give the first analysis of the effectiveness of low-cost detection-based countermeasures against CC-based side-channel attacks. In particular, we analyze two countermeasures - (1) Ciphertext Sanity Check and (2) Decapsulation Failure Check. We demonstrate novel CC-based side-channel attacks that can break both countermeasures efficiently and apply these attacks in the case of Kyber KEM.
 - (a) Our first attack targets the ciphertext sanity check countermeasure, which detects low-entropy ciphertexts used for a large subset of CC-based side-channel attacks. We introduced a novel method to mask malicious low-entropy ciphertexts with the public key to increase their entropy and circumvent the ciphertext sanity check countermeasure. One main advantage of our approach is that, after masking, the attack strategy is the same as the attack strategy without the countermeasure. Our attack works with probability 0.57% and can be restarted if unsuccessful.
 - (b) Our second attack targets the decapsulation failure check countermeasure, which refreshes the secret key upon observing a single decapsulation failure. This countermeasure offers concrete protection as all CC-based attacks rely on multiple invalid/malicious ciphertexts for key recovery. In this work, we propose the first CC-based side-channel attack that only relies on valid ciphertexts to counter the decapsulation failure check countermeasure.
2. We develop a novel solver to obtain the secret key from linear inequalities obtained from leakage. This solver improves upon the Belief Propagation approach, which has been the de-facto approach in prior works [PP21, HPP21, Del22]. Our solver is easier to understand, more efficient to run and can recover the secret key with less than $2\times$ number of linear inequalities compared to the state-of-the-art.

We perform experimental validation of our CC attack using valid ciphertexts, on both the unprotected and masked open-source implementations of Kyber KEM taken from the pqm4 library [KRSS] and mkm4 library [HKL⁺22] respectively. On the unprotected implementations, we can recover the secret key in ≈ 325 traces for the reference implementation and ≈ 7800 traces for the optimized implementation. We also show that our attack can be adapted to both the shuffled and masked implementations in a straightforward manner. While both shuffling and masking increase the attacker’s complexity for key recovery, they do not concretely prevent the attack.

Our work therefore shows that low-cost detection countermeasures can be rendered completely ineffective, and do not offer standalone protection against CC-based side-channel attacks. While these countermeasures are attractive for designers, our work encourages more study towards the development and analysis of new detection-based countermeasures against CC-based side-channel attacks.

2 Preliminaries

2.1 Notations

For any prime q , we let \mathbb{Z}_q denote the field of integers modulo q . Let R_q be the polynomial ring $\mathbb{Z}_q[X]/(X^n + 1)$, then R_q^k is a module of dimension k . Polynomials $a \in R_q$ are denoted using lower case letters. Vectors $\mathbf{a} \in R_q^k$ and matrices $\mathbf{A} \in R_q^{k \times k}$ are denoted in bold using lower and upper case, respectively. When $\mathbf{u}, \mathbf{v} \in R_q^k$, we let $\langle \mathbf{u}, \mathbf{v} \rangle \in R_q$ denote their dot product. The i th coefficient of a polynomial a is denoted as $a[i]$. For any set X , we write $x \stackrel{\$}{\leftarrow} \mathcal{U}(X)$ to denote that x is chosen uniformly at random from X . We denote as χ the Centered Binomial Distribution (CBD) with range $[-\eta, \eta]$. In this case, we abuse notation and let $\mathbf{a} \stackrel{\$}{\leftarrow} \chi(R_q^k)$ mean that each coefficient of each polynomial of $\mathbf{a} \in R_q^k$ is drawn according to χ . Rounding a coefficient of a polynomial $a \in R_q$ from modulus q to modulus p is denoted as $\lfloor a \rfloor_{p \rightarrow q}$.

Let `poly_to_vec` be the function that, given a polynomial in $a \in R_q$, returns the vector in \mathbb{Z}_q^n consisting of the n coefficients of a . Furthermore, let `negashifti` be the function that returns a negacyclic shift of a vector by i positions. That is, if $a = a_0 + a_1x + \dots + a_{n-1}x^{n-1}$, then `negashifti(a)` = $[a_i, \dots, a_0, -a_{n-1}, \dots, -a_{i+1}]$. Using this notation, we can express the coefficients of the product of polynomials a and b in the negacyclic ring R_q in vector form as

$$\text{poly_to_vec}(ab)[i] = \langle \text{negashift}_i(a), \text{poly_to_vec}(b) \rangle.$$

If we extend these two notation for vectors of polynomials $\mathbf{u}, \mathbf{v} \in R_q^k$ by applying the `poly_to_vec` and `negashift` to each of the k polynomials, we get the analogous property

$$\text{poly_to_vec}(\langle \mathbf{u}, \mathbf{v} \rangle)[i] = \langle \text{negashift}_i(\mathbf{u}), \text{poly_to_vec}(\mathbf{v}) \rangle. \quad (1)$$

2.2 Kyber KEM

Kyber is a CCA-secure KEM based on the hardness of the Module-Learning With Errors problem (MLWE) [ABD⁺20]. It offers parameter sets for three NIST security levels 1, 3, and 5, named Kyber512, Kyber768, and Kyber1024 respectively. All instantiations operate over the same anti-cyclic polynomial ring $R_q = \mathbb{Z}_q/(X^n + 1)$ with a prime modulus $q = 3329$ and degree $n = 256$. This ring is used to build the module R_q^k , where $k = 2, 3$ or 4 , for security levels 1, 3, and 5. The CCA-secure Kyber contains, at its core, an IND-CPA secure PKE scheme, which is reviewed next.

<p>Algorithm 1: CPAPKE.KEYGEN.</p> <ol style="list-style-type: none"> 1 $\mathbf{A} \xleftarrow{\\$} \mathcal{U}(R_q^{k \times k})$ 2 $(\mathbf{s}, \mathbf{e}) \xleftarrow{\\$} \chi(R_q^k) \times \chi(R_q^k)$ 3 $\mathbf{b} \leftarrow \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$ 4 return $pk := (\mathbf{A}, \mathbf{b}), sk := \mathbf{s}$ 	<p>Algorithm 2: CPAPKE.ENC.</p> <p>Input: $pk = (\mathbf{A}, \mathbf{b})$ Input: $m \in \mathcal{M}$ Input: $r \xleftarrow{\\$} \{0, 1\}^{256}$</p> <ol style="list-style-type: none"> 1 $(\mathbf{r}, \mathbf{e}_1, \mathbf{e}_2) \xleftarrow{r} \chi(R_q^k) \times \chi(R_q^k) \times \chi(R_q)$ 2 $\mathbf{u} \leftarrow \mathbf{A}^\top \cdot \mathbf{r} + \mathbf{e}_1$ 3 $v \leftarrow \langle \mathbf{b}, \mathbf{r} \rangle + e_2 + \lfloor \frac{q}{2} \rfloor \cdot m$ 4 $\mathbf{c}_1 \leftarrow \lfloor \mathbf{u} \rfloor_{q \rightarrow p}$ 5 $c_2 \leftarrow \lfloor v \rfloor_{q \rightarrow T}$ 6 return $c := (\mathbf{c}_1, c_2)$
<p>Algorithm 3: CPAPKE.DEC.</p> <p>Input: $sk = \mathbf{s}$ Input: $c = (\mathbf{c}_1, c_2)$</p> <ol style="list-style-type: none"> 1 $\mathbf{u}' \leftarrow \lfloor \mathbf{c}_1 \rfloor_{p \rightarrow q}$ 2 $v' \leftarrow \lfloor c_2 \rfloor_{T \rightarrow q}$ 3 $m' \leftarrow v' - \langle \mathbf{s}, \mathbf{u}' \rangle$ 4 $m \leftarrow \lfloor m' \rfloor_{q \rightarrow 2}$ 5 return m 	

2.2.1 IND-CPA Secure Kyber PKE

The IND-CPA secure Kyber PKE scheme consists of three procedures: CPAPKE.KeyGen, CPAPKE.Encrypt and CPAPKE.Decrypt, which are shown in Algorithms 1, 2 and 3, respectively.

Key generation. The key generation procedure shown in Algorithm 3 essentially involves the creation of an LWE instance. The coefficients of the secret key $\mathbf{s} \in R_q^k$ and error $\mathbf{e} \in R_q^k$ are sampled from the narrow centered binomial distribution χ , while coefficients of \mathbf{A} are sampled from the uniform distribution \mathcal{U} . The MLWE instance is computed as $\mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$. The public key is the pair (\mathbf{A}, \mathbf{b}) while the secret key is \mathbf{s} .

Encryption. The encryption procedure shown in Algorithm 2 first samples $\mathbf{r} \in (R_q^k)$, $\mathbf{e}_1 \in R_q^k$ and $e_2 \in R_q$ from χ , which, together with the public key $pk = (\mathbf{A}, \mathbf{b})$, are used to compute two LWE components: $\mathbf{u} \in R_q^k$ and $v \in R_q$. Polynomial v is computed using the 256 bit message $m \in \{0, 1\}^{256}$, for which it is encoded to a corresponding polynomial in R_q by simply multiplying each bit of the message by $\lfloor q/2 \rfloor$. The vector (\mathbf{u}, v) undergo coefficient-wise compression after which they serve as the ciphertext outputs $(\mathbf{c}_1, c_2) \in R_q^k \times R_q$.

Decryption The decryption procedure shown in Algorithm 3 (CPAPKE.Dec) involves the computation of the noisy message polynomial $m' = (v' - \mathbf{s}^T \cdot \mathbf{u}') \in R_q$, where $\mathbf{u}' \in R_q^k$ and $v' \in R_q$ are decompressed versions of values (\mathbf{u}, v) computed during encryption.

The compression and decompression introduces a certain error in \mathbf{u} and v , which we denote as $\Delta \mathbf{u} = \mathbf{u}' - \mathbf{u}$ and $\Delta v = v' - v$, respectively. The noisy message polynomial m' is then decoded to the correct message $m \in \{0, 1\}^{256}$ by simply decoding each of its coefficients $m'[i]$ to 0 or to 1, depending if $m'[i]$ is closer to 0 or to $q/2$.

Let us review why the decryption works. Notice that the noisy message polynomial m' can be represented as follows

$$m' = v' - \langle \mathbf{s}, \mathbf{u}' \rangle = (v + \Delta v) - \langle \mathbf{s}, \mathbf{u} + \Delta \mathbf{u} \rangle.$$

After expanding each term and simplifying it, we get that $m' = m \lfloor \frac{q}{2} \rfloor + \Delta m$, where

$$\Delta m = \langle \mathbf{e}, \mathbf{r} \rangle - \langle \mathbf{s}, \mathbf{e}_1 + \Delta \mathbf{u} \rangle + e_2 + \Delta v. \quad (2)$$

Notice that Δm is made of components sampled from the narrow distribution χ , and therefore m' can be seen as a noisy version of the original message polynomial $\lfloor \frac{q}{2} \rfloor \cdot m$. Kyber security parameters are responsible for ensuring that the coefficients in Δm are small enough so that decryption errors occur only with negligible probability.

Algorithm 4: CCAKEM.KEYGEN.

```

1  $z \xleftarrow{\$} \{0, 1\}^{256}$ 
2  $(pk, sk') = \text{CPA.KEYGEN}()$ 
3  $sk = (sk' || pk || H(pk) || z)$ 
4 return  $pk, sk$ 

```

Algorithm 5: CCAKEM.ENCAPS.

Input: Public key of CCAKEM pk

```

1  $m \xleftarrow{\$} \{0, 1\}^{256}$ 
2  $m \leftarrow H(m)$ 
3  $(\tilde{K}, r) = G(m || H(pk))$ 
4  $c = \text{CPA.ENC}(pk, m, r)$ 
5  $K = \text{KDF}(\tilde{K} || H(c))$ 
6 return  $c, K$ 

```

Algorithm 6: CCAKEM.DECAPS.

Input: Ciphertext of CCAKEM c
Input: Secret key of CCAKEM sk

```

1 Extract  $(sk' || pk || H(pk) || z)$  from  $sk$ 
2  $m' = \text{CPA.DEC}(sk', c)$ 
3  $(\tilde{K}', r') = G(m' || H(pk))$ 
4  $c' = \text{CPA.ENC}(pk, m', r')$ 
5 if  $c = c'$  then
6 |  $K = \text{KDF}(\tilde{K}' || H(c))$ 
7 else
8 |  $K = \text{KDF}(z || H(c))$ 
9 return  $K$ 

```

2.3 IND-CCA Secure Kyber KEM

The IND-CPA secure Kyber KEM can be transformed into IND-CCA secure KEM using the well-known Fujisaki Okamoto (FO) transform [FO99]. This involves instantiation of the CPAPKE.Enc, CPAPKE.Dec procedures of Kyber PKE scheme, along with several instances of hash functions resulting in IND-CCA secure encapsulation (CCAKEYM.Encaps) and decapsulation (CCAKEYM.Decaps) procedures respectively. Algorithm 4 - 6 supplies the details. The main idea is that the randomness required for CPAPKE.Enc is made explicit through a seed r derived from the message m and public key pk . This enables the decapsulation procedure to retrieve the message m from ciphertext ct , compute the seed r , and re-compute the ciphertext ct' .

Subsequently, the recomputed ciphertext ct' is compared with the received ciphertext ct , and this comparison only succeeds for valid ciphertexts and fails for invalid ciphertexts with a very high probability. In theory, the FO transform helps check the validity of ciphertexts through a re-encryption procedure after decryption. Thus, with a very high probability, the attacker only sees decapsulation failures for invalid ciphertexts. This provides strong *theoretical* security guarantees against chosen-ciphertext attacks.

2.4 Prior Works

The decapsulation procedure of Kyber KEM has been subjected to a wide variety of SCA for key recovery [BDH⁺21, RRCB20, HHP⁺21, RRD⁺23]. These attacks can be broadly classified into two categories, based on the knowledge/control over the ciphertexts manipulated by the decapsulation procedure: (1) Known Ciphertext (KC) based SCA and (2) Chosen Ciphertext (CC) based SCA.

2.4.1 Known Ciphertext (KC) based SCA

Known ciphertext attacks only assume knowledge of the ciphertext to the target’s decapsulation procedure. In this respect, the polynomial multiplication operation between the ciphertext component $\mathbf{u} \in R_q^k$ and $\mathbf{s} \in R_q^k$ has been the sole target for side-channel attacks, as it directly manipulates the secret key [PPM17, PP19, YRZ⁺23]. The INTT operation used in the polynomial multiplication operation was targeted by single trace attacks in [PPM17, PP19], while the pointwise multiplication operation was targeted by single trace template style attacks [YRZ⁺23, BBB⁺23] as well as multiple trace Correlation Power Analysis style attacks [MWK⁺22].

Since the polynomial multiplication operation directly manipulates the secret key sk , it becomes an obvious target for such side-channel attacks, and thus the designers can

lay special focus on incorporating extensive countermeasures to protect the polynomial multiplication operation.

2.4.2 Chosen Ciphertext (CC) based SCA

Chosen ciphertext attacks assume the attacker can query the decapsulation procedure with chosen ciphertexts [BDH⁺21, RRCB20, RRD⁺23, UXT⁺22]. The modus operandi of the attack is as follows: The attacker crafts specially structured malicious ciphertexts such that their corresponding decrypted message m (unknown to the attacker) becomes related to the secret key sk (or a part of it). For valid ciphertexts, the message m is independent of the secret key, however for chosen-ciphertexts, the message m is not a key-dependent variable. This ensures that leakage from the entire FO transform can be utilized as an *oracle* to recover the secret key. We can broadly classify existing attacks in this category into three types: (1) Plaintext Checking (PC) Oracle-based SCA [RRCB20, UXT⁺22, TUX⁺23] (2) Full-Decryption (FD) Oracle-Based SCA [XPRO20] and Decryption-Failure (DF) Oracle-Based SCA [BDH⁺21].

2.4.3 Comparing KC-based CCA and CC-based SCA

We observe that KC-based attacks observe leakage for valid ciphertexts, and thus can only exploit leakage from the decryption procedure for key recovery. Moreover, reported attacks have only exploited leakage from the polynomial multiplication operation in the decryption procedure [PPM17, PP19, YRZ⁺23, MWK⁺22]. However, CC-based side-channel attacks form the largest category of attacks on Kyber KEM, and are typically more powerful than KC attacks, as the attacker can exploit leakage from the entire decapsulation procedure for key recovery. Thus, protection against CC-based attacks requires countermeasures for the entire decapsulation procedure, and is therefore naturally costlier compared to protecting KC-based attacks against key recovery attacks. Thus, CC-based side-channel attacks remain the main focus of our work.

There exist concrete proposals for masking as well as combination of masking and shuffling countermeasures to protect against such CC-based attacks. However, they are typically expensive incurring several factors of overhead, particularly in terms of runtime [BGR⁺21, HKL⁺22]. Thus, an alternative line of approach towards countering such attacks has been towards development of low-cost countermeasures that attempt to detect such malicious ciphertexts [XPRO20, RCDB22]. If detected as malicious, the DUT can choose to refresh the secret key, ensuring further exposure for attacks is prevented.

2.5 Detection Countermeasures against CC based SCA

In this section, we broadly discuss two concrete proposals for such detection based countermeasures against CC-based side-channel attacks.

Ciphertext Sanity Check: The malicious ciphertexts used for PC oracle and FD oracle-based attacks are typically very sparse and skewed, with most of the coefficients having a value of 0 [XPRO20, RCDB22]. However, valid ciphertexts are LWE instances whose coefficients are uniformly distributed in the range $[0, q]$. Thus, [RCDB22] proposed to test the entropy of input ciphertexts to detect PC oracle or FD oracle-based attacks. This countermeasure works in the following manner.

In the ciphertext sanity check, the mean and variance is used as measures for the entropy of the input ciphertext. We denote the mean and standard deviation of the coefficients of a polynomial $x \in R_q$ as $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ respectively. The main idea of the countermeasure is then to compute $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ for all the polynomials in the input ciphertext (i.e., \mathbf{c}_1 and \mathbf{c}_2) and reject the ciphertext if the variance of these inputs is too low. One main advantage

of this approach is that the malicious ciphertexts are rejected even before manipulated by the decapsulation procedure. Thus, the attacker cannot observe any leakage corresponding to these malicious ciphertexts thereby offering concrete protection against PC oracle and FD oracle attacks that utilize skewed ciphertexts. Ravi *et al.* [RCDB22] showed that this countermeasure only incurs about 5% overhead in runtime. However, we remark that this countermeasure does not protect against the DF oracle based attacks [BDH⁺21], which already utilize uniformly random ciphertexts that are invalid.

Decapsulation Failure Check: A basic observation of all CC-based side-channel attacks is that all these attacks utilize invalid ciphertexts, that always induce a decapsulation failure. Thus, a more simpler countermeasure at the protocol level, would be to simply refresh the secret key immediately upon observing a decapsulation failure. In this case the attacker is restricted to recovering the secret key with only a single trace. This countermeasure therefore offers concrete protection against all types of CC based side-channel attacks, since most if not all CC based side-channel attacks require multiple traces (few tens to few thousands) for key recovery. In this work, we perform a concrete analysis of both these countermeasures and demonstrate novel CC based side-channel attacks that are capable of breaking both these countermeasures.

3 Analysis of Ciphertext Sanity Check Countermeasure

In this section, we will show that the ciphertext sanity check countermeasure can be bypassed at low cost and with only a small reduction in the accuracy of the attack. This is done by masking a malicious attack ciphertext to look uniformly random using the public key.

Attack Intuition: For the intuitive explanation, we will assume that the public key and ciphertext vectors only consist of one polynomial (i.e., $k = 1$) and that no ciphertext compression is used.

The main idea behind the masking can be intuitively explained as follows: imagine a public key $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$, and an attack ciphertext $(\mathbf{u}_{atk}, v_{atk})$ that would not pass the ciphertext sanity check. The goal is to obtain the message polynomial:

$$\Delta = v_{atk} - \mathbf{u}_{atk} \cdot \mathbf{s}. \quad (3)$$

To this end we adapt our attack ciphertext to $(\mathbf{u}_{atk} + \mathbf{A}, v_{atk} + \mathbf{b})$. The attack ciphertext looks uniformly random and is accepted. The message in this case is:

$$\Delta = v - \mathbf{u} \cdot \mathbf{s} \quad (4)$$

$$= v_{atk} + \mathbf{b} - (\mathbf{u}_{atk} + \mathbf{A}) \cdot \mathbf{s} \quad (5)$$

$$= v_{atk} + \mathbf{A} \cdot \mathbf{s} + \mathbf{e} - \mathbf{u}_{atk} \cdot \mathbf{s} - \mathbf{A} \cdot \mathbf{s} \quad (6)$$

$$= v_{atk} - \mathbf{u}_{atk} \cdot \mathbf{s} + \mathbf{e}. \quad (7)$$

which is the original equation up to the small error e .

One countermeasure against this attack might be to check if the ciphertext is close to the public key, or more specifically whether $(\mathbf{u} - \mathbf{A}, v - \mathbf{b})$ is small. However, such a countermeasure can easily be circumvented by instead of adding the public key, adding a small multiple of the public key (e.g. $(-\mathbf{A}, -\mathbf{b})$ or $(2\mathbf{A}, 2\mathbf{b})$) or a rotated version of the public key (e.g. $(X \cdot \mathbf{A}, X \cdot \mathbf{b})$).

Table 1: Attack parameters used for the CC-based side channel attack using skewed ciphertexts in [RRD⁺23].

$s[0][0] =$		-2	-1	0	1	2
\mathbf{u}_{atk}	v_{atk}	m				
208	624	1	0	0	0	0
208	832	1	1	0	0	0
208	1040	1	1	1	0	0
208	1248	1	1	1	1	0

Attack on Kyber: We now go into more detail on the practical implementation of our attack, where we show how to mask the ciphertexts used for the PC oracle-based side-channel attack. As in previous works, the goal of our attack is to input a chosen ciphertext $(\mathbf{u}_{atk}, v_{atk})$ of the form $\mathbf{u}_{atk} = [y, 0, 0]$ and $v_{atk} = z$ with $|y| < q/4$. The first message bit $m[0]$ is then calculated as:

$$m[0] = \left\lfloor \frac{2}{q} \cdot (z - \mathbf{s}[0][0] \cdot y) \right\rfloor \quad (8)$$

where $m[0]$ is therefore only dependent on the first coefficients of the secret (i.e.) $\mathbf{s}[0][0]$, while all other message bits are zero. The message thus only has two possible values, and the attacker can instantiate a practical PC oracle through side-channel leakage from the decapsulation procedure. The attacker chooses similar values for the tuple (y, z) such that the message can be used as a binary distinguisher for the first secret coefficient. By repeating this procedure for other values of y and z ; and other coefficients of \mathbf{s} , one can retrieve the full secret key. This is a well-known attack procedure as performed in [RRCB20, RRD⁺23]. In [RRD⁺23] the following values of y and z were selected to attack Kyber768:

The attack strategy described above does not pass the ciphertext sanity check, and thus we will have to add the public key to mask the input ciphertext. A first practical problem is that the size of $\mathbf{u} \in R_q^{k \times 1}$ is not equal the size of $\mathbf{A} \in R_q^{k \times k}$, and the same holds for the values $v \in R_q^{1 \times 1}$ and $\mathbf{b} \in R_q^{k \times 1}$. This can be solved by selecting only one row of the matrix \mathbf{A} , which we will indicate with \mathbf{a}^* , and the corresponding polynomial element of the vector \mathbf{b} , which will be indicated with b^* . The introduced error will be indicated e^* and is the corresponding element of the vector \mathbf{e}

A second inconvenience is that the ciphertext is compressed, which means we cannot exactly input the ciphertext. This introduces an error to the values of $\mathbf{a}^* + \mathbf{u}_{atk}$ and $b^* + v_{atk}$ that can be inputted in the attack ciphertext. One can calculate that this results in a message Δ defined as:

$$\Delta = v_{atk} - \mathbf{u}_{atk} \cdot \mathbf{s} + E, \quad \text{where } E = \mathbf{e} - \mathbf{s}^T \cdot \Delta \mathbf{u} + \Delta v.$$

This is essentially the noisy message equation from 2.2.1 where $\mathbf{r} = 1$, $\mathbf{e}_1 = 0$ and $e_2 = 0$ (interestingly, one can see the original unmasked attack as an instance where $\mathbf{r} = 0$). This means that the message Δ calculated in the attack (see 7), has an additional error term due to the compression and decompression.

For power-of-two moduli schemes (e.g., Frodo, Saber), there is typically an additional public key compression of the \mathbf{b} term. The compression, in this case, works as follows: first LSBs are discarded from $\log_2(q)$ bits to $\log_2(p)$ bits in public key compression, then from $\log_2(p)$ to $\log_2(T)$ in ciphertext compression. This means that the public key compression is essentially contained in the ciphertext compression, so its effect can be ignored.

For a given mask (\mathbf{a}^*, b^*) , the value of E is nearly identical for the recovery of different coefficients. This means that each mask (\mathbf{a}^*, b^*) is tied to a certain unknown approximate error E^* . If this approximate error is too large, the attack will fail with high probability, but if this approximate error is low enough, the attack proceeds without the masking

having any effect other than defeating the detection countermeasure. Note that one cannot predict the value of this error E^* in advance. If the key recovery fails, one can restart the attack with a different mask (e.g. using a different row of \mathbf{A} and \mathbf{b} , using a rotated version $(X \cdot \mathbf{a}^*, X \cdot \mathbf{b}^*)$ or using a multiple of the mask $(c \cdot \mathbf{a}^*, c \cdot \mathbf{b}^*)$).

In our experiments on Kyber768, each mask $(\mathbf{a}^*, \mathbf{b}^*)$ had a $p = 57.8\%$ success probability for full key recovery. This means that one has to restart only 0.73 times on average (i.e. $\sum_{i=1}^{\infty} i \cdot p \cdot (1-p)^i$) before hitting a working mask and thus performing a successful attack. We did perform software simulations of our attack, but since the oracle outputs are exactly the same as the PC oracle-based attack in [RRD⁺23], we did not perform side-channel evaluation, as the realization of the oracle has already been demonstrated in [RRD⁺23].

Result: We introduced a method to mask malicious input ciphertext with the public key to avoid the ciphertext sanity check countermeasure. The power of our methodology is that we do not have to change the attack strategy to circumvent the countermeasure: any attack that is possible without the sanity check countermeasure is also possible with the sanity check countermeasure with our masking technique. Our technique works successfully in 57.8% of cases and if the attack is unsuccessful, one can restart the attack with a new mask. This means that to successfully retrieve the secret key, one needs to restart the attack only 0.73 times on average. As an example, we can use the same attack methodology and parameters as in [RRD⁺23], or our technique can be used to mask ciphertexts for other CC based side-channel attacks including FD oracle-based attacks [XPRO20].

4 Analysis of Decapsulation Failure Check Countermeasure

The goal of this section is to develop attacks that can circumvent countermeasures, that refresh the secret key, even upon observing a single decapsulation failure. Defeating this countermeasure calls for a completely new approach, that utilizes valid chosen ciphertexts for key recovery. Our approach stands contrary to the typical state-of-the-art attacks that rely on invalid ciphertexts that fail the re-encryption with high probability.

4.1 Targets for CC-based SCA with Valid Ciphertexts

Since the attacker needs to rely on valid ciphertexts, he/she is limited to targeting operations in the decryption procedure for key recovery. A first possible target would be leakage from the multiplication operation of the secret key $\mathbf{s} \in R_q^k$ with the ciphertext $\mathbf{u} \in R_q^k$. Choosing this operation is quite natural since it directly manipulates the secret key. Previous attacks that perform this are [PPM17, PP19, MWK⁺22], but these attacks can just work with known ciphertexts for key recovery. Moreover, given the direct manipulation of the secret key, the designer might incorporate additional countermeasures to deter key recovery attacks.

This raises the question whether there are other potential operations that can be targeted within the decryption procedure. In this respect, we observe that side-channel leakage from the noisy message polynomial m' can also be used to attack and retrieve the secret key. Exploiting this leakage for practical attacks is not trivial due to two main reasons. Firstly, it is not possible to mount CPA-style attacks that use a divide-and-conquer approach since every coefficient of m' depends on all coefficients of the secret key. Secondly, while the value of m' depends on the secret key, operations manipulating m' do not directly involve the secret key. This means that one will only obtain indirect information on the secret, and no direct information on individual coefficients as in typical CC attacks. Consequently, to obtain the key from the leakage of m' , one needs an additional key recovery step. In this work, we demonstrate how to overcome these obstacles and show

that an attacker who queries only valid chosen ciphertexts can still exploit leakage of m' for efficient key recovery attacks.

4.2 Intuitive Explanation of our Attack

We represent the noisy message polynomial m' as $m' = \lfloor q/2 \rfloor \cdot m + \Delta m$ where Δm is the decryption noise component. We recall from Section 2.2.1, that Δm is linearly dependent on the secret \mathbf{s} and error \mathbf{e} as $\Delta m = \langle \mathbf{e}, \mathbf{e} \rangle - \langle \mathbf{s}, \mathbf{e}_1 + \Delta \mathbf{u} \rangle + e_2 + \Delta v$.

For a valid chosen ciphertext, the attacker knows \mathbf{r}, \mathbf{e}_1 and e_2 , as well as the extra noise factors $\Delta \mathbf{u}$ and Δv , which are caused by the compression and decompression of ciphertexts. Knowledge of $\Delta m \in R_q$ for $2 \cdot k$ ciphertexts results in direct recovery of the secret key \mathbf{s} using linear algebra. The same applies if single coefficients of Δn can be recovered across $2 \cdot k \cdot n$ ciphertexts. As we are limited to a KC attack, we cannot retrieve exact information about Δm . However, in this work, we show that leakage of the coefficients of m' can be used to obtain information on the Hamming weight of Δm , which can then in turn be used to find the sign of the coefficients of m' . From there we can use a linear inequality solver to retrieve the secret vectors \mathbf{s} and error \mathbf{e} .

As we are limited to a KC attack, we cannot retrieve exact information about Δm . However, in this work, we show that leakage of the Hamming weight of coefficients of m' can be used to define bounds on Δm . From there we can use solvers for systems of linear inequalities to retrieve the secret vectors \mathbf{s} and error \mathbf{e} . Our chosen ciphertext attack proceeds in three steps:

Step 1. Recovering $\text{HW}(m')$ through Side-Channels: Manipulation of the coefficients of m' leaks its Hamming Weight (HW), and we assume that the attacker can recover HW of the single coefficients of m' . In Section 5, we will experimentally demonstrate recovery of HW of the coefficients of m' from different implementations of Kyber KEM.

Step 2. Relating $\text{HW}(m')$ with Δm : We observe that for a message bit $m_i = 0$, the distribution of $\text{HW}(m'[i])$ has a very clear bias, that clearly indicates the sign of $\Delta m[i]$. In particular, when $\text{HW}(m'[i])$ is very low, then $\Delta m[i] > 0$ with a very high probability, while a very high value for $\text{HW}(m'[i])$, signifies $\Delta m[i] < 0$ with a very high probability. This information can be used to construct a linear inequality in the coefficients of \mathbf{s} and \mathbf{e} . Generating a sufficient number of inequalities can be used to recover the secret key.

Step 3. Solving Linear Inequalities in \mathbf{s} and \mathbf{e} for Key Recovery: In this step we solve the linear inequalities from Step 2 to retrieve the secret vectors \mathbf{s} and \mathbf{e} . This approach is shown in Section 6. Pessl and Prokop [PP21], Hermelink *et al.* [HPP21] and later Delvaux [Del22] proposed efficient methods based on the Belief Propagation techniques to solve these linear inequalities in the context of fault attacks. We propose a novel greedy solver approach that is simpler to understand, uses lesser compute resources and at the same time requires $2\times$ fewer inequalities to extract the secret key.

4.3 Adversary model

The attacker attempts to recover the long-term secret key sk used by the target's decapsulation procedure of Kyber KEM. We assume physical access to DUT performing decapsulation for power/EM measurements. The attacker has ability to communicate with the target decapsulation procedure with chosen ciphertexts of his/her choice. Note that all used chosen ciphertexts are valid ciphertexts. The attacker has access to a clone device, on which the attacker can control the secret key, in order to capture side-channel measurements for building side-channel templates, and also obtain access to any intermediate variable within the decapsulation procedure.

5 Recovering $\text{HW}(m')$ through Side-Channels

We first present details of our experimental setup, followed by side-channel analysis to detect leakage of single coefficients of m' in different implementations of Kyber KEM.

5.1 Experimental Setup

We performed experiments on reference and optimized implementations of Kyber KEM from the well-known *pqm4* library [KRSS] running on the STM32F407 microcontroller clocked at 24 MHz. We obtain EM side-channel leakage using a Langer RF-U 5-2 near-field EM probe, and the traces were collected using a Lecroy 610Zi oscilloscope at a sampling rate of 1.25 GSam/sec, amplified 30dB with a pre-amplifier and filtered using a 48 MHz low-pass filter.

We split our analysis into two categories: (1) Reference Implementation [ABD⁺20] (2) Assembly-Optimized Implementation [HZZ⁺22].

5.2 SCA of Reference Implementation

Refer to Fig.1 for the assembly code snippet of the subtraction operation between $\mathbf{s}^T \cdot \mathbf{u}$ and v (Line 4 in CPAPKE.Dec procedure in Alg.3), when compiled with the 03 optimization level. The coefficients of m' are computed and stored in memory, sequentially one coefficient at a time, which allows to observe leakage from multiple coefficients simultaneously. To illustrate this, we simultaneously profile leakage of the first 8 coefficients $m'[i]$ for $i \in \{0, 7\}$.

We build 100k valid ciphertexts for random messages m , but choose $m_i = 0$ for $i \in \{0, 7\}$. We perform CPA for the coefficients $m'[i]$ for $i \in \{0, 7\}$ and refer to Fig. 2(a) that shows 8 CPA plots, one for each coefficient $m'[i]$ for $i \in \{0, 7\}$ for the 03 optimized implementation. The peaks for each of the coefficients are easy to distinguish and sufficiently separated from one another indicating clear leakage from individual coefficients.

5.3 SCA of Assembly Optimized Implementation

We studied the assembly-optimized implementation of Kyber KEM, based on the work of Huang *et al.* [HZZ⁺22], and refer to Fig.3 for the assembly code snippet of the assembly-optimized polynomial subtraction operation. As can be seen, it is implemented using highly optimized and vectorized assembly instructions and in particular, we observe that 10 coefficients of the operands are simultaneously loaded into 5 registers using the vectorized `ldmia` instruction (Lines 3, 5).

Thus, each register is packed with two coefficients, where one coefficient is in the higher 16 bits and the other coefficient in the lower 16 bits of the same register. Subsequently,

```

1
2 /* Load Coeff. of A (2 bytes) in r2 */
3 ldrh.w r2, [r4, #2]!
4 /* Load Coeff. of B (2 bytes) in r1 */
5 ldrh.w r1, [r3, #2]!
6 /* Check Loop Counter for Iteration */
7 cmp    r4, r5
8 /* Subtract A from B and store in r2 */
9 sub.w  r2, r2, r1
10 /* Store result in r2 into memory */
11 strh.w r2, [r0, #2]!
12 /* Branch if not end of loop */
13 bne.n  0x8005d30

```

Figure 1: Assembly code snippet of PolySub operation (polynomial subtraction) when compiled with 03 optimized implementation. The target operation that leaks the HW of the coefficients are highlighted in red.

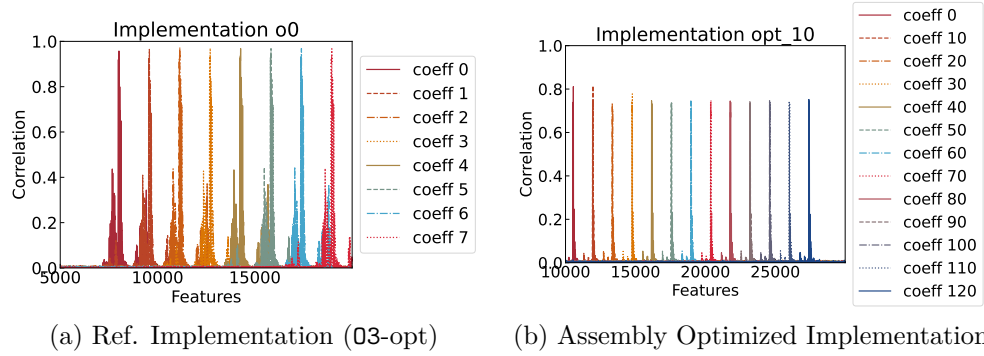


Figure 2: Correlation Power Analysis (CPA) plot corresponding to coefficients of the noisy message polynomial m' for the reference implementation 00-optimized (a), 03-optimized optimization (b) and Assembly optimized implementation (c)

the `usub16` instruction is used to simultaneously subtract two coefficients, and five such instructions are used to subtract 10 coefficients (Line 7-11). Then, the vectorized store instruction (`stmia`) is used to simultaneously store 10 coefficients in memory. We therefore observe that one can only observe simultaneous leakage from 10 coefficients, and leakage for single coefficients can only be observed with a significant amount of noise from the other coefficients that are also stored in parallel.

We will now show leakage detection for the first coefficient $m'[0]$, for which we build 100k ciphertexts for random messages such that, the first message bit $m_i = 0$, while the subsequent nine (9) message bits have a value of 1 (i.e.) $m_i = 1$ for $i \in \{1, 9\}$, and the remaining bits of m' are random. We choose this structure for m , to ensure uniform noise from 9 coefficients $m'[i]$ for $i \in \{1, 9\}$ on $m'[0]$. The same can be repeated by splitting the message m into a set of 10 bits each, where the first bit in each set is 0, while all the other bits are 1.

Refer to Figure 2(c) that shows 13 CPA plots, corresponding to leakage from 13 coefficients (i.e.) $m'[i \cdot 10]$ for $i \in [0, 12]$. Though 10 coefficients are simultaneously stored in memory, we are able to observe leakage from a single targeted coefficient, by setting the other 9 message bits to a value of 1. We can observe the CPA peaks corresponding to the targeted coefficients, and that the peaks are identical and uniformly separated from one another. This demonstrates that simultaneous leakage from one every ten coefficients can

```

1
2 /* Load 10 Coeffs. of A in r3 - r7 */
3 ldmia r1!, {r3, r4, r5, r6, r7}
4 /* Load 5 Coeffs. of B in r8 - r12 */
5 ldmia.w r2!, {r8, r9, r10, r11, r12}
6 /* Simultaneously subtract two coeffs in r3-r8 from r8-r12 */
7 usub16 r3, r3, r8
8 usub16 r4, r4, r9
9 usub16 r5, r5, r10
10 usub16 r6, r6, r11
11 usub16 r7, r7, r12
12 stmia r0!, r3, r4, r5, r6, r7
13 /* Update loop counter */
14 subs.w lr, lr, #1
15 /* Branch if not end of loop */
16 bne.w 0x8005a14

```

Figure 3: Assembly code snippet of PolySub operation (polynomial subtraction) for the highly optimized assembly implementation of Kyber KEM for the ARM Cortex-M4 microcontroller. The target store operation that leaks the HW of the coefficients are highlighted in red.

be exploited for key recovery. In the following, we demonstrate how the detected leakage can be exploited to recover $\text{HW}(m'[i])$.

5.4 Building a Classifier for $\text{HW}(m'[i])$

The HW classifier is built in two phases: (1) Profiling phase and (2) Recovery phase. We illustrate the technique for a single coefficient $m'[0]$ and the same can be repeated for other coefficients.

5.4.1 Profiling Phase

This phase is carried out using leakage from a clone device. From the CPA plot corresponding to $m'[0]$ (Fig.2), we select those features whose correlation value is above a certain threshold Th_0 , as our points of interest (PoI) denoted as \mathcal{P}_0 . We stress that Th_0 is a parameter of the experimental setup, and can be experimentally determined. We use the selected features \mathcal{P}_0 to build a reduced trace set $\mathcal{RT}_{(0,i)}$ for every possible value of $\text{HW}(m'[0]) = i$ for $i \in \{0, 16\}$. We can compute the mean and co-variance matrix of each reduced trace set $\mathcal{RT}_{(0,i)}$ for $i \in \{0, 16\}$, which we denote as $\mu_{(0,i)} \in \mathbb{R}^{\|\mathcal{P}_0\|}$ and $\Sigma_{(0,i)} \in \mathbb{R}^{(\|\mathcal{P}_0\|) \times (\|\mathcal{P}_0\|)}$. Thus, the reduced template for $\text{HW}(m'[0]) = i$ is denoted as $tmp_{(0,i)} = (\mu_{(0,i)}, \Sigma_{(0,i)})$ for $i = \{0, 16\}$. The same procedure can be repeated to build side-channel templates for multiple coefficients of m' .

To perform classification of the HW classes, we have decided to use Random Forest, the reason being is that it has been shown to be successful in the previous works [CDCG22] and it does not have to deal with a more complex hyperparameter tuning usually encountered in more complex models. Random Forest or RF [Bre01] is an ensemble learning algorithm, based on the construction of multiple decision trees. The predictions from these decision trees are combined (for example, through majority voting) in order to achieve a better prediction. The individual decision tree by itself is usually sensitive to small changes in the training data, and when grown large enough, will usually tend to overfit the training data. RF is designed to address these problems of instability in the decision tree, and multiple trees are allowed to grow large, without the need of post-processing. These decision trees are trained on different subsets of the training data using bootstrapping methods (the training data is sampled uniformly with replacement), and the decision is then made by taking the average or majority voting of the decisions given by the trees.

5.4.2 Recovery Phase

In the recovery phase, we obtain a trace tr for a chosen ciphertext from the decryption procedure, and build a reduced trace tr'_0 , corresponding to \mathcal{P}_0 . We utilize the trained RF model for prediction, with the number of trees fixed to $1.5k$, where we observed a good accuracy performance on the validation set. To check if the data imbalance affects the performance, we also verify other statistical tests, such as precision and recall, and obtain a similar score as the accuracy. We then fit in the attack traces to obtain the prediction, which will then be forwarded to the solver to recover the key.

5.4.3 Experimental Results for HW Recovery

We performed experimental validation to recover the HW of single coefficients of m' for both the reference implementation and assembly optimized implementations of Kyber KEM. For the reference implementation compiled with O3 optimization level, we obtain an accuracy of 91.1% in recovering the HW of single targeted coefficients. In case of the assembly optimized implementation, we obtain a much lower accuracy of 32.0% for recovering HW of single targeted coefficients. This much lower accuracy for the HW

classifier can be attributed to the noise due to simultaneous storage of multiple coefficients. However, in the following section, we show that such imperfect HW classifiers with low accuracy do not deter key recovery using valid chosen ciphertexts.

6 Key recovery

We have shown the attacker’s ability to recover $\text{HW}(m')$ (i.e.) HW of individual coefficients of m' . In this section, we demonstrate a key recovery attack exploiting this information. We start by explaining how $\text{HW}(m')$ can be used to derive information about key-dependent noise component Δm and thereby construct linear inequalities in the secret key coefficients. Subsequently, we explain our novel solver to perform full key recovery.

6.1 Recovering Information on Δm using $\text{HW}(m')$

In this section, we consider the type of information we can learn about the decryption noise Δm from the leaked Hamming weight of the message $\text{HW}(m')$. The intuitive idea of our approach is to characterize the relation between this Hamming weight and the sign of Δm . From this characterization, we can assign to each Hamming weight a probability of the sign being positive or negative.

6.1.1 A First Approach

It is known that the coefficients of m' are distributed as a narrow bell-shaped discrete distribution around $q/2$ or 0 in the following manner

$$m'[i] = \begin{cases} q/2 \pm \Delta m[i] & \text{for } m_i = 1, \\ 0 \pm \Delta m[i] & \text{for } m_i = 0, \end{cases} \quad (9)$$

where the standard deviation of the noise $\Delta m[i]$ is $\sigma(\Delta m[i]) \approx 79$.

In most practical implementations of Kyber, such as *pqm4*, coefficients of m' are represented as signed 16-bit integers. Then, because of the narrow distribution of the noise $\Delta m[i]$, we hypothesized that it should be possible to learn something about the size of $\Delta m[i]$ from the observation of $\text{HW}(m'[i])$.

Figure 4 shows the distribution of $\text{HW}(m'[i])$ when $\Delta m[i] \geq 0$ and $\Delta m[i] < 0$, considering $m[i] = 0$ and $m[i] = 1$. We can see that, when $m[i] = 0$, there is a clear distinction between positive and negative errors. This happens since a small negative error when added to a zero message bit (i.e. $m[i] \lfloor q/2 \rfloor = 0$), results in a small negative number, which has a very high Hamming weight in the typically used two-complement representation. Similarly, small positive errors would result in values around 0 on the positive side, which in turn have very small Hamming weights. On the other hand, when $m[i] = 1$, then the HW distributions for positive and negative noise $\Delta m[i]$ are very close to each other, and therefore it is hard to obtain meaningful information on $\Delta m[i]$ from $\text{HW}(m'[i])$.

From this observation, if an attacker learns the HW of multiple $m'[i]$ where $m[i] = 0$, they can build a series of inequalities on $\Delta m[i]$. These types of inequalities can then be solved using Belief Propagation algorithms [PP21, Del22, HPP21]. While this approach is theoretically sound, it wastes some of the information that comes from the Hamming weight. In the following section, we discuss how to obtain more information than only the sign of $\Delta m[i]$ from $\text{HW}(m'[i])$.

6.1.2 Obtaining Tighter Inequalities

The Hamming weight of the message can not only be used to infer information about the sign of Δm but also on the possible values that Δm can take. That is, a certain

Table 2: Maximum and minimum values of $m'[i]$ for each possible value of its Hamming weight that was observed in 100,000 decryptions using Kyber’s reference implementation.

HW ($m'[i]$)	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
max $m'[i]$	0	256	320	328	338	351	324	319	316	251	253	-32	-16	-8	-4	-2	-1
min $m'[i]$	0	1	-257	-256	-254	-250	-352	-357	-314	-338	-319	-332	-324	-305	-321	-257	-1
Spread	0	255	577	584	592	601	676	676	630	589	572	586	308	297	317	255	0

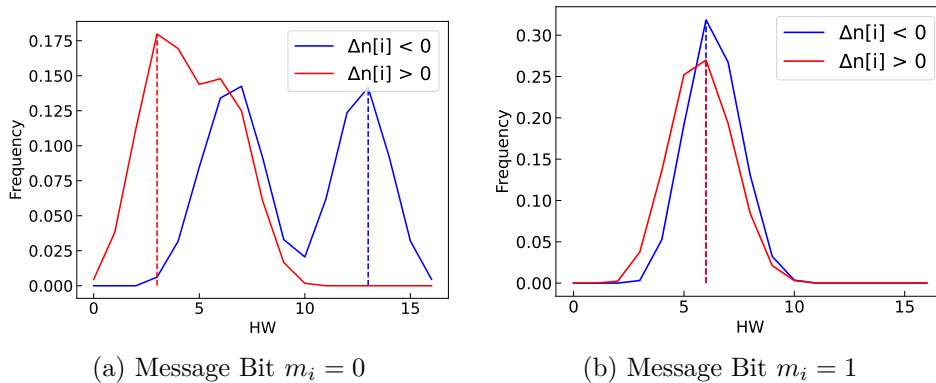
Hamming weight can only be linked to a limited number of possible values of Δm and some Hamming weights can only occur within a limited range of values. Therefore, it is possible to compute both a lower bound and an upper bound on $m'[i]$ based on the Hamming weight. There is one important caveat when implementing this idea in practice: efficient Kyber implementations perform lazy modular reduction, and therefore it is not guaranteed that $-q/4 \leq m'[i] < q/4$ before the reduction. While it is not easy to give a theoretical estimate on the bounds on $m'[i]$ for each Hamming weight from 0 to 16, we can run a simulation to compute these values.

Table 2 shows the maximum and minimum values of $m'[i]$ observed for each possible Hamming weight, considering Kyber’s reference implementation. The tightness of the inequalities are represented by the spread, which is simply the difference between $\max m'[i] - \min m'[i]$. It is also possible to see the effect of the lazy reduction. Notice how one would expect $\max m'[i]$ to be -64 for the case when $\text{HW}(m'[i]) = 10$, since it’s two’s complement representation `0b1111111111000000` makes it the maximum value that can be the result of a narrow distribution around 0 with the given Hamming weight. However, instead we observe 253 because $\text{HW}(253 + q) = \text{HW}(0b110111111110) = 10$.

Now, for each observed $\text{HW}(m'[i])$ with $m[i] = 0$, an attacker is able to construct two inequalities, one for the maximum and one for the minimum. This gives more information to the inequalities solver, which we will cover next.

6.2 Solving Inequalities to Find The Secret Key

In this section, we formalize how a system of linear inequalities related to the secret key is built with SCA information. We then construct a new solver to retrieve the secret key from these inequalities.

Figure 4: Distribution of HW of message polynomial coefficients $m'[i]$ for positive error ($\Delta n[i] > 0$) and negative error ($\Delta n[i] < 0$), when $m_i = 0$ (a) and $m_i = 1$ (b)

6.2.1 Defining the System of Linear Inequalities Related to the Key

In the previous section, we saw that when $m[i] = 0$, then $m'[i] = \Delta m[i]$, and $\text{HW}(m'[i])$ gives us information on $\Delta m[i]$. Suppose an attacker asks the target device to decrypt a number κ of ciphertexts, whose corresponding messages m_1, \dots, m_κ . Similarly, let $\mathbf{r}_j, (\mathbf{e}_1)_j, \mathbf{u}_j, (e_2)_j$ and v_j be the values resulting from the encryption of the corresponding ciphertexts, for $j = 1$ to κ , which the attacker knows.

Now, for any indexes $(j, i) \in \{1, \dots, \kappa\} \times \{0, \dots, 255\}$, we know that

$$\Delta m[i] = \langle \mathbf{r}_j, \mathbf{e} \rangle [i] - \langle (\mathbf{e}_1)_j + \Delta \mathbf{u}_j, \mathbf{s} \rangle [i] + (e_2)_j [i] + \Delta v_j [i].$$

But, from Equation 1, we know that

$$\begin{cases} \langle \mathbf{r}_j, \mathbf{e} \rangle [i] = \langle \text{negashift}_i(\mathbf{r}_j), \text{poly_to_vec}(\mathbf{e}) \rangle, \\ \langle (\mathbf{e}_1)_j + \Delta \mathbf{u}_j, \mathbf{s} \rangle [i] = \langle \text{negashift}_i((\mathbf{e}_1)_j + \Delta \mathbf{u}_j), \text{poly_to_vec}(\mathbf{s}) \rangle. \end{cases}$$

Therefore, if we let

$$\mathbf{h}_{(i,j)} = \left[\text{negashift}_i(\mathbf{r}_j) \parallel -\text{negashift}_i((\mathbf{e}_1)_j + \Delta \mathbf{u}_j) \right] \in \mathbb{Z}_q^{2kn},$$

then

$$\langle \mathbf{h}_{(i,j)}, [\text{poly_to_vec}(\mathbf{e}) \parallel \text{poly_to_vec}(\mathbf{s})] \rangle + (e_2)_j [i] + \Delta v_j [i] = \Delta m_j [i].$$

Now, when $m_j[i] = 0$, then the attacker can use Table 2 to obtain two values $\omega_{(i,j)}$ and $\Omega_{(i,j)}$ such that $\omega_{(i,j)} \leq \Delta m_j [i] \leq \Omega_{(i,j)}$. This means the attacker can build two inequalities in this case

$$\begin{cases} \mathbf{h}_{(i,j)} [\text{poly_to_vec}(\mathbf{e}) \parallel \text{poly_to_vec}(\mathbf{s})] + (e_2)_j [i] + \Delta v_j [i] - \omega_{(i,j)} & \geq 0, \\ -\mathbf{h}_{(i,j)} [\text{poly_to_vec}(\mathbf{e}) \parallel \text{poly_to_vec}(\mathbf{s})] - (e_2)_j [i] - \Delta v_j [i] + \Omega_{(i,j)} & \geq 0. \end{cases} \quad (10)$$

The attacker can then select a number γ of pairs (j, i) such that $m_j[i] = 0$ and get a linear system of inequalities. This system of inequalities is represented by a matrix $\mathbf{H} \in \mathbb{Z}_q^{2\gamma \times 2kn}$, and a vector $\mathbf{w} \in \mathbb{Z}_q^{2\gamma}$ as defined next, based Equation 10.

Algorithm 7: Greedy key recovery algorithm.

```

Input:  $\mathbf{H} \in \mathbb{Z}_q^{2\gamma \times 2kn}$ ,  $\mathbf{w} \in \mathbb{Z}_q^{2\gamma}$ 
1  $\mathbf{x} \leftarrow \mathbf{0} \in \mathbb{Z}_q^{2kn}$ 
2 for  $\text{it} = 1$  to  $\text{max\_iterations}$  do
3    $\text{action\_scores} \leftarrow \{\}$   $\triangleright$  dictionary where keys are actions and values are scores
4   for  $(j, v) \in \{0, \dots, 2kn - 1\} \times \{-\eta, \dots, \eta\}$  do
5      $\text{action\_scores}[j, v] \leftarrow \text{COMPUTESCORE}(\mathbf{x}, (j, v), \mathbf{H}, \mathbf{w})$ 
6    $\text{best\_actions} \leftarrow$  List of keys  $(j, v)$  with the  $\alpha_{\text{it}}$  highest scores in  $\text{action\_scores}$ 
7   for  $(j, v)$  in  $\text{best\_actions}$  do
8      $\mathbf{x}[j] \leftarrow \mathbf{x}[j] + v$   $\triangleright$  Apply action  $(j, v)$  to  $\mathbf{x}$ 
9 return  $\mathbf{x}$ 

```

Each pair of rows of matrix \mathbf{H} consists of the vectors $\mathbf{h}_{(i,j)}$ and $-\mathbf{h}_{(i,j)}$. Entries of \mathbf{w} consist of the factors $((e_2)_j [i] + \Delta v_j [i] - \omega_{(i,j)})$ and $(-(e_2)_j [i] - \Delta v_j [i] + \Omega_{(i,j)})$. Then, if γ is sufficiently large and if the inequalities are sufficiently tight, then a small solution \mathbf{x} to the linear system $\mathbf{H}\mathbf{x} + \mathbf{w} \geq \mathbf{0}$ should give us the secret $[\text{poly_to_vec}(\mathbf{e}) \parallel \text{poly_to_vec}(\mathbf{s})]$.

One may be tempted to use every single inequality that is generated. However, some of them will be very loose. There is an important trade-off when selecting a good spread

Algorithm 8: Score function.

<p>Input: $\mathbf{x} \in \mathbb{Z}^{2kn}$, $(j, v) \in \{0, \dots, 2kn - 1\} \times \{-\eta, \dots, \eta\}$, $\mathbf{H} \in \mathbb{Z}^{\gamma \times 2kn}$, $\mathbf{w} \in \mathbb{Z}^{\gamma}$</p> <pre> 1 score \leftarrow 0 2 $\mathbf{H}_j \leftarrow$ jth column of \mathbf{H} 3 $\mathbf{t} \leftarrow \mathbf{H}\mathbf{x} + \mathbf{w} + v\mathbf{H}_j$ ▷ Resulting target vector when action (j, v) is applied 4 for $i = 1$ to γ do ▷ For each inequality, lower score by the distance it is from being satisfied 5 if $\mathbf{t}[i] < 0$ then 6 score \leftarrow score $- \mathbf{t}[i]$ 7 return score </pre>

value to filter pairs of inequalities. If it is too small, then a larger number of traces may be needed, but if it is too large, then the algorithm must be prepared to deal with a potentially very large number of inequalities. To find the spread minimizing the number of ciphertexts needed for a successful attack, we ran a simple simulation considering all spread values in Table 2, and obtained the value of 317.

6.2.2 New Solver for Key Recovery

Until now, the most efficient solver available for recovering the key from inequalities was the Belief Propagation algorithm proposed by Pessl and Prokop [PP21] with the improvements by Delvaux [Del22]. This solver, however, has an important limitation: it is not efficient when dealing with lots of inequalities and requires a large amount of RAM.

We propose a simpler greedy algorithm to find the key from the inequalities constructed from the Hamming weight of the noisy message coefficients. Our algorithm is presented as Algorithm 7. It starts with $\mathbf{x} = \mathbf{0} \in \mathbb{Z}^{2kn}$ and then, on each iteration, it performs a number of actions that, according to the heuristic score function in Algorithm 8, should guide \mathbf{x} closer to the solution. Each action is a pair (j, v) , where j is an index of \mathbf{x} and v is an integer that is added to $x[j]$. To obtain a faster convergence, we let the number α_{it} of actions to be applied on iteration it to start as a large number and decay exponentially in the number of iterations.

The score function for actions defined in Algorithm 8 works as follows. First, it applies the action on \mathbf{x} and verifies which inequalities are not satisfied after the action was applied. Then, for each unsatisfied inequality, it penalizes the action score with the numerical distance from the offending number to the closest value that satisfies the inequality.

Figure 5 shows a comparison between our greedy search solver and the one based on belief propagation (BP) [Del22]. We can see that, despite its simplicity, our solver compares favorably with respect to the number of ciphertexts needed to recover the key. Furthermore, it appears to be more robust with respect to noise in the SCA observations.

Notice that the performance of the belief propagation algorithm degrades fast. Therefore, we could not generate more points for larger standard deviations, because, since we need to run a binary search for the computing the minimum number of ciphertexts needed, the simulation took too long to finish. For comparison, when the SCA noise standard deviation is $\sigma = 0.5$, the BP algorithm takes, on average, more than 10 minutes, while our algorithm takes less than 20 seconds.

6.3 Experimental Results for Key Recovery

We implemented the entire attack on Kyber768 (recommended parameter set) to recover 10 random secret keys, and will separately explain the results for the reference (O3 optimized) as well as assembly optimized implementation of Kyber KEM.

Reference Implementation: We remark that any randomly generated message m contains an average of 128 zero bits, and thus the attacker can simultaneously exploit

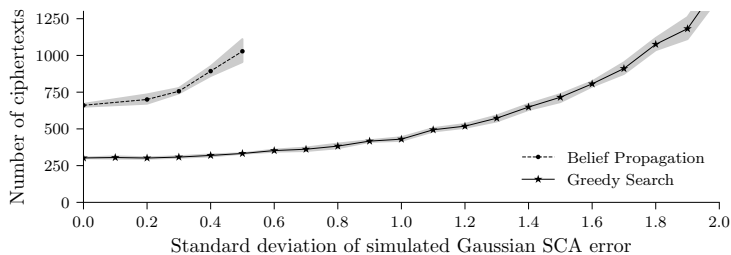


Figure 5: Comparison of our Greedy Search Solver compared to Belief Propagation for key recovery, considering different levels of Gaussian noise (i.e.) $\sigma = 0.0$ to 2.0

leakage from all the 128 bits. However, to simplify our attack demonstration and setup, we only exploit leakage from the first 8 zero bits of the message. Using this leakage, we were able to recover all the 10 secret keys in about ≈ 5200 traces. Thus, exploiting leakage from 128 zero bits in a message means that full key recovery is possible in just 325 traces, which correlates well with our attack simulations as presented in Figure 5.

Optimized Implementation: We showed in Section 5.3 that simultaneously storing 10 coefficients in memory ensures the attacker can only exploit leakage from one every 10 coefficient of m' . For illustration, we exploited leakage from 13 coefficients of m' (i.e.) m_{10*i} for $i \in \{0, 12\}$, and were able to recover all secret keys in about ≈ 7800 traces. Thus, the number of traces is $24\times$ higher than that to attack the reference implementation (325). This can be attributed to the reduced number of exploited coefficients and the reduced accuracy of the HW classifier ($\approx 32\%$ for the optimized implementation, compared to $\approx 91\%$ for the reference implementation). In the optimized implementation, we exploited leakage of single coefficients from 13 vectorized store instructions. However, we remark that an attacker can exploit from several more vectorized store instructions, out of a total of 26 vectorized instructions for key recovery.

Thus, it is clear that leakage of m' can be efficiently exploited by an attacker for key recovery with valid chosen ciphertexts. In the following, we demonstrate the applicability of our attack to the shuffling and masking countermeasures.

7 Attacking Shuffled Implementation

We now consider the case when the operations manipulating the coefficients of m' are shuffled. In this case, the attacker does not know when the zero message bits are being processed, and therefore even if they know that $m[i] = 0$, they cannot tell if $m'[i]$ have a high or low Hamming weight. But notice that, since the leakage of the message polynomial coefficients still exists, the attacker can recover HW of single coefficients of m' in the same manner as an unprotected implementation. This can be done by using templates created using leakage from a clone device, on which the attacker has knowledge of the random permutation used for every execution. However, we cannot directly use the same key recovery procedure without knowing the permutation used for processing these coefficients on the target device.

Attack Idea: Since the attacker can only exploit leakage from coefficients $m'[i]$ corresponding to $m_i = 0$, we propose to choose a message m that has only a very small number of zero bits. We remark that the attacker can explicitly choose the value of the message for valid chosen ciphertexts. Suppose that the chosen message m has $\theta = 2$ null bits at positions i and j , and that the attacker sees a sequence of Hamming weights $W = (w_1, w_2, \dots, w_n)$ being processed. Then, if this sequence contains two very small values $w_{i^*}, w_{j^*} \leq 1$ or two

very large values $w_{i^*}, w_{j^*} \geq 15$, then i^* and j^* must be associated with the zero positions i and j in some way. Once this is done, we can proceed using the same key recovery strategy presented in the previous section, considering either 1 or 15 as the HW for both $m'[i]$ and $mt[j]$, depending on the HW values that were observed.

The main difficulty in using this idea in practice is that we need really tight intervals on a Hamming weight w to associate it to a zero message bit $m[i] = 0$, since we need to be sure to exclude all possible valid HW associated with $m[i] = 1$. This results in the observations within these intervals being very rare. Unfortunately, using $\theta = 2$ as in the example does only give us $\binom{n}{2} = 32640$ different ciphertexts, which are not enough for us to make a high number of such rare observations. However, we can relax this idea and use messages with $\theta = 4$ null entries, but accept inequalities in case we see 2 extreme values. Then all indexes corresponding to $m[i] = 0$ be counted as having an extreme Hamming weight. Notice that this would generate some invalid inequalities.

As a proof of concept, we tested this idea and obtained a successful recovery, assuming perfect classification of the HW of the coefficients $m'[i]$. With 300000 inequalities, out of which around 16% were wrong, we were able to recover the key using a little more than 8 million ciphertexts. Although the number of ciphertexts is very large, it does support the validity of the approach, and we leave improving this value for future work.

8 Attacking Masked Implementation

In the masked implementation of the decryption procedure, the message polynomial m' is additively masked, and thus the attacker can observe leakage from the d independent arithmetic shares for each coefficient $m'[i]$ of the message polynomial, where $d - 1$ is the order of masking. We performed our experiments on the first order masked implementation of Kyber768 from the open-source `mkm4` library [HKL⁺22].

Consider the `MaskedPolySub` operation responsible for the masked computation of the subtraction $m' = v - \langle \mathbf{s}, \mathbf{u} \rangle$. Let $w = \langle \mathbf{s}, \mathbf{u} \rangle \in R_q$, and consider its arithmetic shares $w = w_0 + w_1 + \dots + w_{d-1}$. The computation is done in two steps. First, it uses the `poly_sub` to compute $v - w_0$. Then the other shares are simply negated. Notice that `poly_sub` is similar to the unprotected version.

While the `mkm4` library utilizes an assembly optimized implementation for the polynomial subtraction operation, for our analysis, we modified this routine with a C based implementation and compiled it with the `O0` compiler optimization level for simplicity. Our analysis of the assembly code also reveals that the attacker can observe leakage of the individual coefficients of all the arithmetic shares of m' . We performed CPA to detect leakage from coefficients of all the arithmetic shares of m' and refer to Figure 6

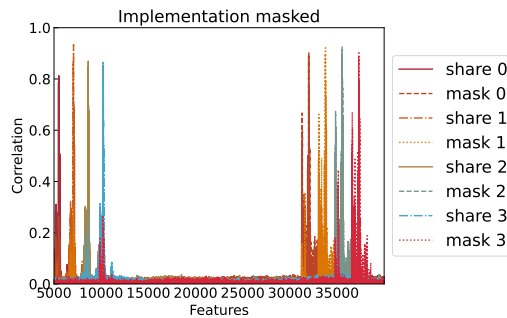


Figure 6: Correlation Power Analysis (CPA) plot for four coefficients of the two arithmetic shares of m' , from the reference implementation of masked Kyber KEM, compiled with `O0` optimization level.

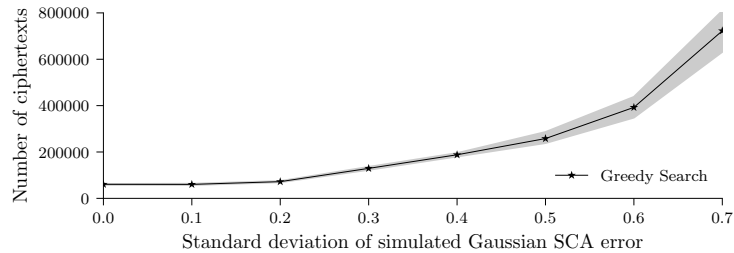


Figure 7: Number of ciphertexts required for key recovery considering different levels of Gaussian noise

for the CPA plots, where we can observe leakage from coefficients of all shares of m' . For the classification of HW, we used the same technique that was used for the unprotected implementation and we were able to obtain an average accuracy of 94% for the HW classifier of all the targeted coefficients. We now show how the HW of coefficients of arithmetic shares of m' can be plugged into our novel greedy search algorithm for key recovery.

Key Recovery. From the key recovery point of view, the main difference when attacking the masked implementation is the construction of inequalities. Remember that, in the unprotected case, we used tables of minimums and maximum values of $m'[i]$ for each possible HW. This idea can be adapted for building inequalities for the masked implementation.

For the first order masked implementation, we build two 17×17 tables. In these tables, rows and columns represent the Hamming weights of the first and second shares of $m'[i]$. The entry with index (ω_0, ω_1) of the first table is the maximum observed value for $m'[i]$ when its shares have Hamming weights ω_0 and ω_1 . The second table is similarly constructed, except that we take the minimum instead of maximum. To build the maximums and minimums tables, we used the code from *mkm4* implementation and observed 100000 decryptions. We remark that no SCA is needed for building this matrix, as they are only dependent on the implementation, not on the device. The resulting inequalities can then be directly plugged into the key recovery algorithm. Figure 7 shows the results performance of our key recovery algorithm in the masked case. Notice that we did not consider the Belief Propagation algorithm here because the required number of inequalities is too large for it to be efficient.

9 Conclusion

We performed the first security analysis of two detection-based SCA countermeasures against CC-based side-channel attacks: (1) Ciphertext Sanity Check and (2) Decapsulation Failure Check, demonstrating practical attacks against both countermeasures. We first report a novel attack to circumvent the ciphertext sanity checking, by simply applying the public key as a mask to a maliciously crafted ciphertext. We circumvent the decapsulation failure check, by proposing the first CC based side-channel attack that relies on valid ciphertexts for key recovery. Our attack exemplarily exploits leakage from the noisy message polynomial in the decryption procedure for full key recovery. We also introduced a simpler and improved inequality solver, that can recover the secret key with less than half the inequalities compared to these previous methods based on Belief Propagation.

We performed experimental validation of our attack on reference and optimized implementations of Kyber KEM on the STM32F4 microcontroller, requiring between $\approx 325-7800$ traces for full key recovery. We show how our attack can be adapted to both the shuffled and masked implementations, with appropriate increase in number of traces for key recovery.

Our work therefore shows that low-cost detection countermeasures can be rendered completely ineffective, and do not offer standalone protection against CC-based side-channel attacks. While these countermeasures are attractive for designers, our work encourages more study towards the development and analysis of new detection-based countermeasures against CC-based side-channel attacks.

Acknowledgment

Part of this work was discussed during the Dagstuhl Seminar-23152 titled "Secure and Efficient Post-Quantum Cryptography in Hardware and Software", Apr.2023 and we are thankful to Schloss Dagstuhl, its staff, organizers and participants of Dagstuhl Seminar-23152.

Jan-Pieter D’Anvers is funded by FWO (Research Foundation - Flanders) as junior postdoctoral fellow (contract number 133185). In addition, this work was supported by CyberSecurity Research Flanders with reference number VR20192203, the European Commission through the DIGITAL-SIMPLE project Be-QCI with contract number 101091625, and the Horizon 2020 research and innovation program Belfort ERC Advanced Grant 101020005.

References

- [AAC⁺22] Gorjan Alagic, Daniel Apon, David Cooper, Quynh Dang, Thinh Dang, John Kelsey, Jacob Lichtinger, Carl Miller, Dustin Moody, Rene Peralta, et al. Status report on the third round of the NIST post-quantum cryptography standardization process. Technical report, National Institute of Standards and Technology, 2022.
- [ABD⁺20] Roberto Avanzi, Joppe W. Bos, Leo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Kyber (version 3.0): Algorithm specifications and supporting documentation (October 1, 2020). *Submission to the NIST post-quantum project*, 2020.
- [AHKS22] Amin Abdulrahman, Vincent Hwang, Matthias J Kannwischer, and Daan Sprenkels. Faster Kyber and Dilithium on the Cortex-M4. *Cryptology ePrint Archive*, 2022.
- [BBB⁺23] Estuardo Alpirez Bock, Gustavo Banegas, Chris Brzuska, Łukasz Chmielewski, Kirthivaasan Puniamurthy, and Milan Šorf. Breaking dpa-protected kyber via the pair-pointwise multiplication. *Cryptology ePrint Archive*, 2023.
- [BDH⁺21] Shivam Bhasin, Jan-Pieter D’Anvers, Daniel Heinz, Thomas Pöppelmann, and Michiel van Beirendonck. Attacking and defending masked polynomial comparison for lattice-based cryptography. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2021(3):334–359, Jul. 2021.
- [BGR⁺21] Joppe Willem Bos, Marc Olivier Gourjon, Joost Renes, Tobias Schneider, and Christine van Vredendaal. Masking kyber: First-and higher-order implementations. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2021(4):173–214, 2021.
- [Bre01] Leo Breiman. Random Forests. *Machine Learning*, 45(1):5–32, 2001.

- [CDCG22] Brice Colombier, Vlad-Florin Drăgoi, Pierre-Louis Cayrel, and Vincent Grosso. Profiled side-channel attack on cryptosystems based on the binary syndrome decoding problem. *IEEE Transactions on Information Forensics and Security*, 17:3407–3420, 2022.
- [Del22] Jeroen Delvaux. Roulette: A diverse family of feasible fault attacks on masked Kyber. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 637–660, 2022.
- [FO99] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In *Annual International Cryptology Conference*, pages 537–554. Springer, 1999.
- [HHP⁺21] Mike Hamburg, Julius Hermelink, Robert Primas, Simona Samardjiska, Thomas Schamberger, Silvan Streit, Emanuele Strieder, and Christine van Vredendaal. Chosen ciphertext k-trace attacks on masked CCA2 secure Kyber. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 88–113, 2021.
- [HKL⁺22] Daniel Heinz, Matthias J Kannwischer, Georg Land, Thomas Pöppelmann, Peter Schwabe, and Daan Sprenkels. First-order masked kyber on arm cortex-m4. *Cryptology ePrint Archive*, 2022.
- [HPP21] Julius Hermelink, Peter Pessl, and Thomas Pöppelmann. Fault-enabled chosen-ciphertext attacks on kyber. In *Progress in Cryptology–INDOCRYPT 2021: 22nd International Conference on Cryptology in India, Jaipur, India, December 12–15, 2021, Proceedings 22*, pages 311–334. Springer, 2021.
- [HZZ⁺22] Junhao Huang, Jipeng Zhang, Haosong Zhao, Zhe Liu, Ray CC Cheung, Çetin Kaya Koç, and Donglong Chen. Improved plantard arithmetic for lattice-based cryptography. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2022(4):614–636, 2022.
- [KRSS] Matthias J. Kannwischer, Joost Rijneveld, Peter Schwabe, and Ko Stoffelen. mupq/pqm4: Post-quantum crypto library for the ARM Cortex-M4. <https://github.com/mupq/pqm4>.
- [KRSS19] Matthias J Kannwischer, Joost Rijneveld, Peter Schwabe, and Ko Stoffelen. pqm4: Testing and benchmarking NIST PQC on ARM Cortex-M4. In *Second PQC Standardization Conference: University of California, Santa Barbara and co-located with Crypto 2019*, pages 1–22, 2019.
- [MWK⁺22] Catinca Mujdei, Lennert Wouters, Angshuman Karmakar, Arthur Beckers, Jose Maria Bermudo Mera, and Ingrid Verbauwhede. Side-channel analysis of lattice-based post-quantum cryptography: Exploiting polynomial multiplication. *ACM Transactions on Embedded Computing Systems*, 2022.
- [PP19] Peter Pessl and Robert Primas. More practical single-trace attacks on the number theoretic transform. In *International Conference on Cryptology and Information Security in Latin America*, pages 130–149. Springer, 2019.
- [PP21] Peter Pessl and Lukas Prokop. Fault attacks on CCA-secure lattice KEMs. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 37–60, 2021.

- [PPM17] Robert Primas, Peter Pessl, and Stefan Mangard. Single-trace side-channel attacks on masked lattice-based encryption. In *International Conference on Cryptographic Hardware and Embedded Systems*, pages 513–533. Springer, 2017.
- [RCDB22] Prasanna Ravi, Anupam Chattopadhyay, Jan Pieter D’Anvers, and Anubhab Baksi. Side-channel and fault-injection attacks over lattice-based post-quantum schemes (kyber, dilithium): Survey and new results. *ACM Transactions on Embedded Computing Systems*, 2022.
- [RPBC20] Prasanna Ravi, Romain Poussier, Shivam Bhasin, and Anupam Chattopadhyay. On configurable SCA countermeasures against single trace attacks for the NTT. In *International Conference on Security, Privacy, and Applied Cryptography Engineering*, pages 123–146. Springer, 2020.
- [RRCB20] Prasanna Ravi, Sujoy Sinha Roy, Anupam Chattopadhyay, and Shivam Bhasin. Generic side-channel attacks on CCA-secure lattice-based PKE and KEMs. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2020(3):307–335, 2020.
- [RRD⁺23] Gokulnath Rajendran, Prasanna Ravi, Jan-Pieter D’Anvers, Shivam Bhasin, and Anupam Chattopadhyay. Pushing the limits of generic side-channel attacks on lwe-based kems-parallel pc oracle attacks on kyber kem and beyond. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 418–446, 2023.
- [TUX⁺23] Yutaro Tanaka, Rei Ueno, Keita Xagawa, Akira Ito, Junko Takahashi, and Naofumi Homma. Multiple-valued plaintext-checking side-channel attacks on post-quantum kems. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 473–503, 2023.
- [UXT⁺22] Rei Ueno, Keita Xagawa, Yutaro Tanaka, Akira Ito, Junko Takahashi, and Naofumi Homma. Curse of re-encryption: A generic power/em analysis on post-quantum kems. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 296–322, 2022.
- [XPRO20] Zhuang Xu, Owen Pemberton, Sujoy Sinha Roy, and David Oswald. Magnifying side-channel leakage of lattice-based cryptosystems with chosen ciphertexts: The case study of Kyber. *IACR Cryptol. ePrint Arch.*, 2020:912, 2020.
- [YRZ⁺23] Bolin Yang, Prasanna Ravi, Fan Zhang, Ao Shen, and Shivam Bhasin. Stamp-single trace attack on m-lwe pointwise multiplication in kyber. *Cryptology ePrint Archive*, 2023.