

Ciphertext malleability in Lattice-Based KEMs as a countermeasure to Side Channel Analysis

Pierre-Augustin Berthet* ^C

Hensoldt France SAS, Plaisir, France

LTCI, Télécom Paris, Palaiseau, France

berthet@telecom-paris.fr

Abstract. Due to developments in quantum computing, classical asymmetric cryptography is at risk of being breached. Consequently, new Post-Quantum Cryptography (PQC) primitives using lattices are studied. Another point of scrutiny is the resistance of these new primitives to Side Channel Analysis (SCA), where an attacker can study physical leakages. In this work, we discuss an SCA vulnerability due to the ciphertext malleability of some PQC primitives exposed in a work by Ravi et al. We propose a novel countermeasure to this vulnerability exploiting the same ciphertext malleability and discuss its practical application to several PQC primitives. We also extend the seminal work of Ravi et al. by detailing their attack on the different security levels of a post-quantum Key Encapsulation Mechanism (KEM), namely FrodoKEM. We also provide a generalisation of their attack to different parameters which could be used in future similar primitives.

Keywords: Post-Quantum Cryptography, Lattice-Based KEM, FrodoKEM, ML-KEM, SABER, Side Channel Analysis, Countermeasure

1. Introduction

With the emergence of quantum computing, new primitives for asymmetric cryptography are studied. The National Institute of Standards and Technology (NIST) launched a competition [1, 12] to select a future Post-Quantum Cryptography (PQC) standard. The first Key Encapsulation Mechanism (KEM) to be standardised is the Module Lattice-based Key Encapsulation Mechanism (ML-KEM) or FIPS 203 [23]. This KEM is based on the CRYSTALS-Kyber candidate [3]. However, other finalists from

*Supported by Agence de l’Innovation de Défense, grant 2022156 Thèse CIFRE Défense

^CCorresponding author

the third round of competition are using similar mathematical problems, such as SABER [4] and FrodoKEM [22].

Another concern for new PQC primitives is their resistance to physical attacks. First introduced by Kocher [19], Side-Channel Analysis (SCA) is a kind of cryptanalysis which targets the physical environment of a cryptosystem and analyses potential leakages (e.g. execution timing, heat, power consumption, electromagnetic emissions) to try and discover correlations between sensitive data and the physical data the attacker observes.

The resistance of PQC primitives to SCA has already been studied. Several works summarise existing papers in terms of attacks and possible countermeasures, such as Ravi et al.[29] and Canto et al.[10] for ML-KEM. Most noticeably, ciphertext malleability, a vulnerability of several NIST PQC lattice-based candidates, is presented by Ravi et al.[28].

1.1. Our contributions

In this work, we present a countermeasure to side-channel opponents targeting the decoding of the message in PQC Lattice-based KEM. The attacker uses ciphertext malleability to do so and, according to Ravi et al. [28], generic countermeasures such as masking [11, 17, 25] or shuffling [15, 35] are not sufficient. The countermeasure uses the same ciphertext property to perform the attack as a defence mechanism. As a consequence, the attacker cannot exploit ciphertext malleability to retrieve the secret.

However, from a high-level perspective, the countermeasure reuses existing primitive functions. Although side-channel attacks against these functions can circumvent the countermeasure, it is fully compatible with other generic countermeasures. We recommend its deployment alongside a shuffling countermeasure to have a message side-channel resistance similar to a masking and shuffling countermeasure but at lower costs. In essence, the countermeasure presented in this work “displaces” the leakage point from a function where generic countermeasures are not effective to one where they are.

The countermeasure initially intervenes in the `PKE.Decrypt` procedure of the KEM. We also discuss extending its application to the remainder of the `KEM.Decaps` procedure. We discuss the specificities of this extension for three NIST PQC competition candidates, namely SABER, FrodoKEM and ML-KEM.

We also extend the seminal attack work from Ravi et al. [28], originally against ML-KEM [23], by providing a more detailed insight on how the attack can be performed against the different parameter settings of the post-quantum algorithm FrodoKEM. Ravi et al. [28] attack works by adding biases to the ciphertext to flip specific bits of the output of the decoding of the message during decapsulation. Assuming 8-bit (resp. 16-bit) storage of the message in memory, we can recover the message of Frodo-640 in 9 (resp. 17) traces. For Frodo-976, the adaptation of Ravi et al. [28] attack requires 10 (resp. 19) traces. We propose more optimal choices of biases than the adaptation of Ravi et al. [28] attack by studying the impact of all possible biases instead of only focussing on the ones that flip bits. Our ciphertext choices only require 8 (resp. 14) traces for Frodo-976. For Frodo-1344, the adaptation requires 9 (resp. 17) traces, and our improved choice of biases only 7 (resp. 13).

The optimal choices of biases in our improvement of Ravi et al.[28] attack were done heuristically by brute-forcing all possibilities. A Jupyter Notebook with our methods is available at the following link:

<https://github.com/Pierre-Augustin-Berthet/decode-lpr>

The paper is structured as follows. In Section 2 we introduce notations and detail the LPR framework and the different primitives that use this framework in the third round of the NIST PQC competition. Ciphertext malleability is presented in Section 3 with the seminal work of Ravi et al. [28] and our further extension of their attack to the FrodoKEM algorithm. A generalisation of the attack is also discussed. The countermeasure is discussed throughout Section 4, with descriptions of its overall strategy, extension to the entirety of KEM.Decaps, scalability, compatibility with other generic countermeasures, as well as a discussion on its impact against other SCAs. Section 5 concludes the paper.

2. Preliminaries

2.1. Notations

We denote by $q \in \mathbb{N}$ a modulus and \mathbb{Z}_q the set $\mathbb{Z}/q\mathbb{Z}$. $n \in \mathbb{N}$ a message length, and R_q the devolution ring $\mathbb{Z}_q[X]/\langle X^n + 1 \rangle$. We denote by \mathcal{B} the set of eight bits (or byte) $\{0, 1\}^8$ or \mathbb{F}_2^8 . We denote by $\mathcal{U}(\mathcal{K})$ the uniform distribution over a set \mathcal{K} .

Remark 2.1. In this paper, we discuss three different primitives, each with their specific and conflicting notation. Their notation will be used for their description; however, we unify the notation for the other sections of the paper.

The Hamming Weight of a binary message m is the number of non-zero bits it contains. We denote it $HW(m)$. We denote the difference between the Hamming Weight of a message m and of a biased message m' by $\Delta = HW(m') - HW(m)$.

We introduce the concept of an X-classes distinguisher. Let \mathcal{I} be an input set for a function F and \mathcal{S} the output set. We have a X-classes distinguisher if we can build X distinct classes of subsets of \mathcal{I} from subsets of \mathcal{S} such as $F(\mathcal{I}) = \mathcal{S}$. Consequently, a class is denoted by $S \in \mathcal{S} \rightarrow I \in \mathcal{I}$, that is, the observation of S implies that the input of F has the shape I . In this paper, \mathcal{I} contains binary words composed of the letters 0 and 1. The letter \cdot is used if the information on a letter of the binary word is not certain. On the other hand, elements of \mathcal{S} are composed of a sign “+” or “-” followed by a positive integer. For example, if we consider

$$F : \mathcal{I} \rightarrow \mathcal{S}, F(00) = +1, F(01) = +0, F(10) = +1, F(11) = -1, \quad (1)$$

we have a two-classes as well as a three-classes distinguishers. The two-classes distinguisher is $\{+1 \rightarrow \cdot 0, (+0, -1) \rightarrow \cdot 1\}$. The three-classes one is $\{+1 \rightarrow \cdot 0, +0 \rightarrow 01, -1 \rightarrow 11\}$.

2.2. Generic SCA countermeasures

To avoid correlations between side channel leakage and sensitive variables, several generic countermeasures have been proposed. They tend to add randomness in one way or another to the computations.

For example, *shuffling* [15, 35] randomises the execution order of the algorithm whenever possible. Although this does not eliminate leakage, leakage detection often requires several attempts and data alignment between those attempts to be successful. Shuffling foils the alignment and thus increases the number of attempts required to successfully exploit the leakage.

The core idea behind *blinding* [21, 31] is to add to the computations a random value that will be eliminated in the later stages of the function. As a consequence, it randomises the side-channel leakage of the function for each of its iterations. Blinding methods tend to be specific to the functions they protect and are less generic.

Finally, *masking* [11, 17, 25] aims to divide the sensitive data into several random *shares*. Those shares are processed by the same algorithm separately and reassembled in the later stages of the protected function to ensure the correctness. This forces a side-channel attacker to use more probing hardware as it must recover each share to breach the implementation. It can be considered the most studied and popular generic countermeasure as it benefits from formal proofs and models. A simple example of masking is the *boolean masking*, which masks bits:

$$\text{For } x \in \{0, 1\} \text{ and } r \leftarrow \mathcal{U}(\{0, 1\}), \text{Mask}(x) = (x \oplus r, r), \text{Unmask}((x \oplus r, r)) = (x \oplus r) \oplus r \quad (2)$$

Remark 2.2. The generic countermeasures presented here all intervene at an algorithmic level. Other generic countermeasures can be set up at lower levels, including directly at a hardware level with *shielding* [24], for example. Shielding is performed by putting a physical barrier between the opponent and the targeted hardware, avoiding any leakage.

2.3. LWE/LWR PKE using lattices

Several NIST PQC candidates rely on the Learning With Error (LWE) problem on lattices and its variant the Learning With Rounding (LWR). They are based on the Luybashevsky-Peikert-Regev (LPR) framework [20], described in Algorithm 1. The error distribution on a set \mathcal{K} for LPR is indicated by $\chi(\mathcal{K})$. It can be deterministically computed from a seed r , in this case it will be denoted by $\chi(r, \mathcal{K})$.

The `Encode` function is a compression function. It corresponds for a parameter d to a mapping of elements of \mathbb{Z}_{2^d} to \mathbb{Z}_q . The function `Decode` performs a mapping of elements of \mathbb{Z}_q to \mathbb{Z}_{2^d} . They use the rounding to the nearest integer denoted by $\lceil \cdot \rceil$. When applied to a matrix or a vector or a polynomial, these functions are applied separately on each coefficient of the matrix, vector, or polynomial. This rule applies to matrix and vector of polynomials, these functions are applied separately on the coefficients of the polynomials. They are defined as follows:

$$\forall \beta \in \mathbb{Z}_{2^d}, \text{Encode}(\beta, d) = \left\lceil \frac{q}{2^d} \cdot \beta \right\rceil \bmod q, \quad (3)$$

$$\forall \alpha \in \mathbb{Z}_q, \text{Decode}(\alpha, d) = \left\lceil \frac{2^d}{q} \cdot \alpha \right\rceil \bmod 2^d. \quad (4)$$

We can also use a sector representation to illustrate the effect of the `Decode` function as seen in Figure 1. The possible inputs are represented as positions on the circle boundary, with values ranging from 0 to $q \equiv 0$ counterclockwise. As the output depends on which interval the input is part of, the interval is drawn as a sector of the circle, and the sector is labelled with its corresponding output. For example, in Figure 1b, if the input of `Decode` is between $\frac{q}{8}$ and $\frac{3q}{8}$, it will be mapped to 01.

Algorithm 1: LPR Encryption Scheme [28, 20]

```

1 Procedure PKE.KeyGen():
2   a  $\leftarrow \mathcal{U}(\mathbb{Z}_q)$ 
3   s, e  $\leftarrow \chi(\mathbb{Z}_q)$ 
4   t = a  $\times$  s + e
5   return  $pk = (\mathbf{a}, \mathbf{t})$ ,  $sk = (s)$ 
6
7 Procedure PKE.Encrypt( $pk, m \in \mathcal{B}^{32}, r \in \mathcal{B}^{32}$ ):
8   s', e', e''  $\leftarrow \chi(r, \mathbb{Z}_q)$ 
9   u = a  $\times$  s' + e'
10  v' = t  $\times$  s' + e''
11  x = Encode( $m, d$ )
12  v = v' + x
13  return  $ct = (\mathbf{u}, \mathbf{v})$ 
14
15 Procedure PKE.Decrypt( $ct, sk$ ):
16  x' = (v - u  $\times$  s)
17   $m' = \text{Decode}(\mathbf{x}', d)$ 
18  return  $m'$ 

```

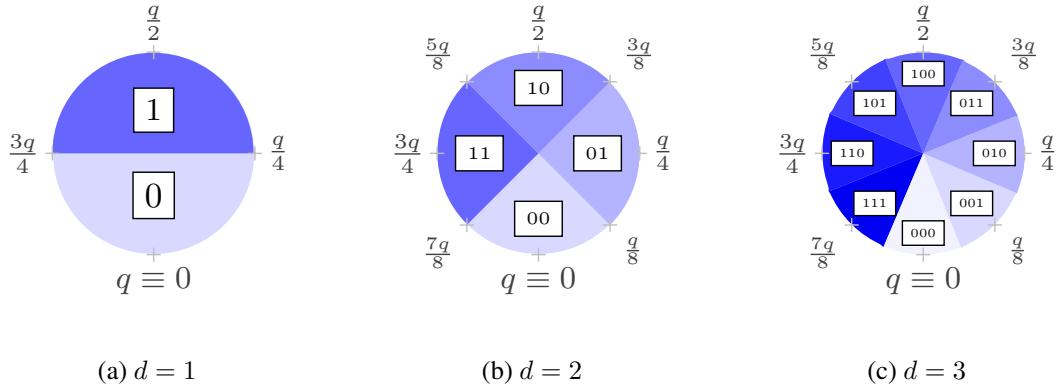


Figure 1: Sector representation of the decoding function for different parameters d , with the intervals on the circle boundary counter-clockwise and the sectors labeled with the corresponding outputs to the intervals

2.4. Fujisaki-Okamoto Transform

The LPR framework is not IND-CCA secure (INDistinguishable-Chosen Ciphertext Attack) as an attacker can modify **u** and **v** to recover information on **s** from m' . To counteract this attack, LPR-

based primitives use an adaptation of the Fujisaki-Okamoto Transform (FOT) [13] as described in Algorithm 2.

Algorithm 2: Fujisaki-Okamoto Transform for LPR [28, 20]

```

1 Procedure KEM.KeyGen():
2    $(pk, s) = \text{PKE.KeyGen}()$ 
3    $z \leftarrow \mathcal{U}(\mathcal{B}^{32})$ 
4   return  $pk, sk = (s, z)$ 
5

6 Procedure KEM.Encaps( $pk$ ):
7    $m \leftarrow \mathcal{U}(\mathcal{B}^{32})$ 
8    $r = \text{PRF}(m, pk)$                                 // Pseudorandom Function
9    $ct = \text{PKE.Encrypt}(pk, m, r)$ 
10   $K = \text{KDF}(r, ct)$                                 // Key Derivation Function
11  return  $ct, K$ 
12

13 Procedure KEM.Decaps( $ct, pk, sk$ ):
14   $m' = \text{PKE.Decrypt}(sk, ct)$ 
15   $r' = \text{PRF}(m', pk)$ 
16   $ct' = \text{PKE.Encrypt}(pk, m', r')$ 
17  if  $ct == ct'$  then
18    | return  $K = \text{KDF}(r', ct')$ 
19  end
20  else
21    | return  $K = \text{KDF}(z, ct')$ 
22  end

```

2.5. FrodoKEM

FrodoKEM [22] is a NIST PQC third-round alternate finalist. Although it was not selected for the fourth round, it is selected by the German Bundesamt für Sicherheit in der Informationstechnik (BSI) and the French Agence Nationale de la Sécurité des Systèmes d’Information (ANSSI) as a more conservative option than ML-KEM. It was noticeably used in the first French diplomatic telegram using post-quantum cryptography on 30 November 2022, and is currently submitted to the International Organisation for Standardisation (ISO). An Internet-Draft (I-D) has also been submitted to the Internet Engineering Task Force (IETF). Thus, it is still relevant to study its resistance to SCAs.

2.5.1. Parameters

Remark 2.3. The notations used in the description of each primitive in this paper are the ones used in their respective specification papers. However, outside of their description, we use the unified notation defined in Section 2.1.

FrodoKEM relies on the LWE problem and the LPR framework. It uses matrices over $\mathbb{Z}_q^{\bar{m} \times n}$, $\mathbb{Z}_q^{\bar{m} \times \bar{n}}$ and $\mathbb{Z}_q^{\bar{n} \times n}$. FrodoKEM offers three levels of security and different sets of parameters for each level:

Table 1: Parameters sets for FrodoKEM

	NIST security level	n	q	$\bar{m} \times \bar{n}$	B
Frodo-640	I	640	2^{15}	8×8	2
Frodo-976	III	976	2^{16}	8×8	3
Frodo-1344	V	1344	2^{16}	8×8	4

The most interesting parameter for this work is the compression factor B applied to the message in FrodoKEM. It is equivalent to the parameter d from Equations 3 and 4. The compression functions are defined as *Frodo.Encode* for *Encode* and *Frodo.Decode* for *Decode*.

Remark 2.4. In FrodoKEM, the parameter n does not denote the message length but rather a matrix size parameter. The message length is given by $l = B \times \bar{m} \times \bar{n}$.

2.6. ML-KEM

ML-KEM [23] is a slight modification of CRYSTALS-Kyber [3], a KEM selected by NIST [1]. ML-KEM is the first post-quantum KEM standard published by the NIST.

2.6.1. Parameters

Its design adapted from LPR relies on several instances of the Module-LWE (M-LWE) and Module-LWR (M-LWR) problems. It uses vectors in R_q^k and square matrices in $R_q^{k \times k}$. ML-KEM offers three levels of security and different sets of parameters for each level:

Table 2: Parameters sets for ML-KEM

	NIST security level	n	q	k	η_1	η_2	d_u	d_v	d_m
ML-KEM-512	I	256	3329	2	3	2	10	4	1
ML-KEM-768	III	256	3329	3	2	2	10	4	1
ML-KEM-1024	V	256	3329	4	2	2	11	5	1

The compression factor applied to the message in ML-KEM is denoted d_m in Table 2 as ML-KEM does not use a specific notation for the parameter d from Equations 4 and 3 when applied to the message, as it is always 1. Another interesting point is the use of compression to reduce the size of the ciphertext in ML-KEM, the relevant parameters being d_u and d_v . The compression functions are defined as *Decompress_q* for *Encode* and *Compress_q* for *Decode*.

2.7. SABER

SABER [4] is a NIST PQC third round finalist. It was not selected for the fourth round nor as a standard by the NIST or the BSI. Its future deployment for real world application is uncertain. Nonethe-

less, its similarities with ML-KEM makes its study relevant. Especially, it relies on a power of 2 modulus rather than a Solinas/Crandall prime modulus like ML-KEM and uses the M-LWR problem rather than M-LWE.

2.7.1. Parameters

SABER offers three levels of security and different sets of parameters for each level:

Table 3: Parameters sets for SABER

	NIST security level	n	q	p	T	l	μ
LightSaber	I	256	2^{13}	2^{10}	2^3	2	10
Saber	III	256	2^{13}	2^{10}	2^4	3	8
FireSaber	V	256	2^{13}	2^{10}	2^6	4	6

As SABER relies only on M-LWR, it uses several hidden compressions which are defined as functions (Decode and Encode) in the specification paper. The compression for the encoding of the message is tied to the parameter p . The parameter T is used for ciphertext compression, a method also used in ML-KEM.

3. Exploiting ciphertext malleability: existing work and generalisation

3.1. Seminal work: Targeted message bit flip for $d=1$

Ravi et al.[28] introduced an attack exploiting what they called “ciphertext malleability” in the LPR framework. They perform a chosen-ciphertext attack with the aim of flipping a bit of the secret message. Given a side-channel Oracle that we will denote by \mathcal{O}_{SCA} , they can recover the exact value of any message bit.

The oracle is defined as follows: When queried, \mathcal{O}_{SCA} provides the Hamming Weight of a register where an output of $\text{PKE}.\text{Decrypt}$ is stored. This can be performed in side-channel against any implementation on a microcontroller which stores the message computed by $\text{PKE}.\text{Decrypt}$ within a group of registers in memory before further computations. Therefore, the attack can be easily performed in a real-world setting. In their work, Ravi et al. [28] consider a perfect Hamming Weight distinguisher as \mathcal{O}_{SCA} . If such a distinguisher is not available, they recommend using an imperfect one several times and to perform a majority vote.

The main metric used in side-channel attacks is the number of traces, i.e. measurements, required to recover the targeted secret. This number depends on many parameters, e.g. the Signal-to-Noise Ratio (SNR) of the measures or the size of the registers in which the secret is stored. For example, the seminal attack by Ravi et al. [28] requires 256 calls to the oracle against MLKEM [23], without the calls to recover the original message Hamming Weight. Taking into account the storage of the message in 8-bit words and an optimal SNR, they parallelise the calls to the oracle and reduce the attack cost to only 9 traces. In this work, we consider the performance of the attack using both traces

and number of queries to the oracle, as some countermeasures like shuffling require to recover the Hamming Weight of the entire message for the attack to still work, and thus remove the possibility of using parallel attacks. The Hamming Weight of the entire message can be recovered by adding the Hamming Weights of its registers.

The ciphertext malleability is due to the function `Decode`. When its input is shifted by adding a specific bias modulus q , this can alter the output in a predictable way. Using \mathcal{O}_{SCA} before and after the chosen ciphertext attack, it is possible to deduce the exact value of any bit of the message. The biases have the form $+\frac{qk}{2^d}$, where $k < 2^d$ an integer and d the compression factor of the function `Decode`. In the case of MLKEM [23], as $d = 1$ and q is not a multiple of 2^d , the bias injected is rounded, i.e. $+\lceil \frac{q}{2} \rceil = +\frac{q+1}{2} = 1665$.

The attack by Ravi et al. [28] against MLKEM is carried out as follows:

1. The Hamming Weight of the original message m is recovered with \mathcal{O}_{SCA} .
2. The bias $+\lceil \frac{q}{2} \rceil$ is injected into the ciphertext to target a specific bit of the message. To target the i^{th} bit of the message, we add the bias to the ciphertext \mathbf{v} during the `PKE.Decrypt` procedure at the start of `KEM.Decaps` in order to force the computation of $\text{Decode}(\mathbf{v} + \lceil \frac{q}{2} \rceil X^i - \mathbf{u} \times \mathbf{s})$.
3. The target device performs its computations and a call to \mathcal{O}_{SCA} retrieves the Hamming Weight of the biased message m' .
4. The results are compared: $\Delta = \text{HW}(m') - \text{HW}(m) = \begin{cases} +1 & \text{if the message bit is 0;} \\ -1 & \text{if the message bit is 1.} \end{cases}$

Remark 3.1. For a parallelised attack, the oracle \mathcal{O}_{SCA} only needs to recover the Hamming Weight of the register which contains the biased part of the message. Hence, the attack can be performed on each register separately in parallel to lower the number of traces.

The representation in sectors allows for a better understanding of how ciphertext malleability can be exploited. Adding a specific bias to the input of `Decode` “rotates” the sectors and alters the output accordingly. An example is given in Figure 2, where one can bit-flip the output of the `Decode` function with parameter $d = 1$ by adding $\frac{q}{2}$ to the input.

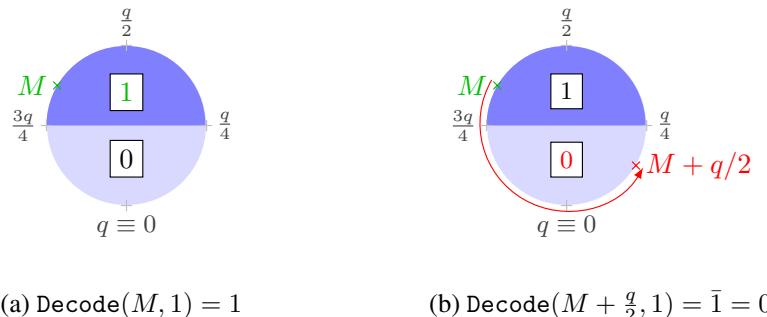


Figure 2: Bitflipping the output of the decoding for parameter $d = 1$

3.2. Generalising targeted message bit flip

Ravi et al. [28] focused on the ML-KEM primitive, where $d = 1$. In this section, we extend their targeted bit flip strategy to any value of the parameter d . A higher value of d also implies that a single call to `Decode` (α, d) outputs several bits of message at a time.

First, we observe that, for any integer $k < 2^d$,

$$\begin{aligned} \text{Decode}\left(\alpha + \frac{qk}{2^d}, d\right) &\equiv \left\lceil \frac{2^d}{q} \left(\alpha + \frac{qk}{2^d} \right) \right\rceil \bmod 2^d \\ &\equiv \left\lceil \frac{2^d}{q} \alpha + k \right\rceil \bmod 2^d \\ &\equiv \left\lceil \frac{2^d}{q} \alpha \right\rceil + k \bmod 2^d \\ &\equiv \text{Decode}(\alpha, d) + k \bmod 2^d \end{aligned}$$

When $d = 1$, this results in XORing the single bit of the output of `Decode` $(\alpha, 1)$ with 1. For higher values of d , k is added to the output. When k is a power of 2, we have the following theorem:

Theorem 3.2. For any value of d , there exists a two-classes distinguisher that gives the value of at least one bit of the output of `Decode` (α, d) with absolute certainty, and this for each bit of the output.

Proof:

Let $k = 2^j$ be the power of 2 for an integer j . Let $m = \text{Decode}(\alpha, d)$ and $m' = \text{Decode}(\alpha, d) + k \bmod 2^d$. Let $\Delta = \text{HW}(m') - \text{HW}(m)$. We use the following binary description: $m = \sum_{i=0}^{d-1} m_i * 2^i$. Adding k to m has the following results:

$$\begin{cases} \text{if } m_j = 0, \text{ then } m'_j = 1 \text{ and } \Delta = +1; \\ \text{if } m_j = 1, \text{ then } m'_j = 0 \text{ but we have a carry to propagate.} \end{cases} \quad (5)$$

If there is a carry to propagate, we have $\Delta = -1 + \epsilon$, with ϵ the impact of carry propagation on the Hamming Weight. The propagation of the carry can be seen as adding $k' = 2k$ to a variant of m , denoted m'' , such as $\forall i \neq j, m_i = m''_i$ and $m''_j = 0 \neq m_j$. Thus, reapplying the results from Equation 5, we have $\epsilon = +1$ and $\Delta = +0$, or $\epsilon = -1 + \epsilon'$, with ϵ' the impact of the propagation of the next carry. Recursively, we have $\epsilon \leq +1$ and $\Delta \leq +0$. We can thus always build the following two-classes distinguisher:

$$\text{For the bias } + \left\lceil \frac{qk}{2^d} \right\rceil, \text{ with } k = 2^j, j \in \mathbb{N}, \{\Delta = +1 \rightarrow m_j = 0, \Delta \leq +0 \rightarrow m_j = 1\}. \quad (6)$$

□

A consequence of Equation 6 is that the efficiency of the Ravi et al.[28] targeted bit flip attack is 1 bit per query for any value of d . Equation 6 guarantees that there is no uncertainty about the

value of the bit m_j . Hence, the targeted bit flip attack requires d queries per instance of the function $\text{Decode}(\alpha, d)$ in the attacked primitive in addition to the queries required to recover the original Hamming Weight of the message.

In a real-world scenario, the message is stored in memory and the queries can be parallelised. As a consequence, Ravi et al. [28] attack can recover the entire message in 9 traces if it is stored in 8-bit registers, 17 traces if it is stored in 16-bit registers, only if d is a power of 2 lesser or equal to the size of the register. Otherwise, the outputs of the function Decode might overlap between two registers, and thus the attack cannot be completely parallelised. For example, this is the case in the second security level of FrodoKEM [22], where $d = 3$.

3.3. The FrodoKEM example

In their work, Ravi et al.[28] detailed their attack for ML-KEM [23] and gave indications on how to extend it to SABER [4] and Frodo-640 [22] in their appendix. In this section, we focus on FrodoKEM [22]. We extend the work of Ravi et al.[28] to cover the higher security parameters of FrodoKEM [22] as some peculiarities remain regarding how to attack them.

Remark 3.3. As the modulus in FrodoKEM [22] is a power of two, we omit the rounding $\lceil \cdot \rceil$ in the biases.

3.3.1. Existing targets

As the attack is based on the Hamming Weight, the preferred target is microcontroller implementation, especially on the Cortex M4 family of microcontrollers. In the case of FrodoKEM [22], we can cite as a reference the *pqm4* [18] Github repository, under the Tag “Round 3”. However, this implementation only contains the first security level of FrodoKEM [22], i.e., Frodo-640. The second level of security of FrodoKEM [22] was first implemented by Howe et al. [16] with improved memory allocation. Further improvements were proposed by Bos et al. [7], using Single Instruction Multiple Data (SIMD) to parallelise some calculations. The first implementation with the third security level, that is, Frodo-1344, is found in a work by Bos et al. [8] with a design centred on memory stack optimisation. Although this allows for the implementation of this third security level of FrodoKEM [22], it comes with a cost in time performance.

Regarding the size of the registers used to store the output of the message decoding, the only publicly available implementation, *pqm4* [18], uses 16-bit registers.

3.3.2. How to perform the attack

We follow the same procedure as the Ravi et al. [28] targeted bit flip attack. In FrodoKEM [22], the input of the function Decode in $\text{PKE}.\text{Decrypt}$ at the beginning of $\text{KEM}.\text{Decaps}$ is a matrix in $\mathbb{Z}_q^{\bar{m} \times \bar{n}}$. The decoding is performed on each coefficient of the matrix separately. To target a specific area of the message, we modify the ciphertext¹ \mathbf{v} , which is itself a matrix in $\mathbb{Z}_q^{\bar{m} \times \bar{n}}$, by adding bias at the correct

¹With FrodoKEM [22] notation: C

coefficient. We then compute $\text{KEM}.\text{Decaps}$ with the biased ciphertext. The attack of Ravi et al. [28] against FrodoKEM [22] is carried out as follows for each coefficient of \mathbf{v} :

1. The Hamming Weight of the unbiased message is recovered with \mathcal{O}_{SCA} ,
2. The bias is injected within the ciphertext \mathbf{v} at the correct coefficient,
3. The Hamming Weight of the biased message is recovered with a call to \mathcal{O}_{SCA} ,
4. Steps 2 and 3 are repeated until enough information is recovered to determine the value of the message bits linked to the targeted coefficient of \mathbf{v} .

Remark 3.4. From an implementation point of view, the elements of \mathbb{Z}_q are encoded in 16 bits. This implies that the most significant bits of these coefficients are always 0 since the modulus of Frodo-640 is lower than 2^{16} . However, the ciphertexts are packed in bytes, where these extra zeros are removed. This implies that injecting bias into the ciphertext \mathbf{v} can impact up to 3 bytes of the packed ciphertext. Unpacking the ciphertext does not alienate the value of the ciphertext coefficient, and thus the value of a biased ciphertext coefficient as well.

For Frodo-976 and Frodo-1344, the packing and unpacking is simply a concatenation of coefficients of the matrix. The bias injection affects 2 bytes of the packed ciphertext.

3.3.3. Frodo-640

The first FrodoKEM security level uses the function Decode with the parameter $d = 2$, denoted B in the FrodoKEM specification [22]. This implies mapping to two bits, and hence $2^d = 4$ possible values. Table 4 summarises the impact on the Hamming Weight of the output of Decode when adding $\frac{q}{4}$, $\frac{q}{2}$ or $\frac{3q}{4}$ to the input compared to the Hamming Weight of the output without any input bias.

Table 4: Impact of biasing the decoding function input with parameter $d = 2$ on the Hamming Weight of the output

Initial input	Initial mapping	$+\frac{q}{4}$	$+\frac{q}{2}$	$+\frac{3q}{4}$
$[-\frac{q}{8}, \frac{q}{8}]$	00	+1	+1	+2
$[\frac{q}{8}, \frac{3q}{8}]$	01	+0	+1	-1
$[\frac{3q}{8}, \frac{5q}{8}]$	10	+1	-1	+0
$[\frac{5q}{8}, \frac{7q}{8}]$	11	-2	-1	-1

We have the following distinguishers for each bias used in Table 4:

- $+\frac{q}{4}$:
 - two-classes: $\{+1 \rightarrow .0, (+0, -2) \rightarrow .1\}$
 - three-classes: $\{+1 \rightarrow .0, +0 \rightarrow 01, -2 \rightarrow 11\}$

- $+\frac{q}{2}$: two-classes: $\{+1 \rightarrow 0, -1 \rightarrow 1\}$
- $+\frac{3q}{4}$:
 - two-classes: $\{-1 \rightarrow 1, (+0, +2) \rightarrow 0\}$
 - three-classes: $\{-1 \rightarrow 1, +0 \rightarrow 10, +2 \rightarrow 00\}$

Ravi et al. [28] targeted bit flip attack can be performed in two queries with biases $+\frac{q}{2}$ and $+\frac{q}{4}$. Given the lists of two-classes distinguishers and Table 4, the two-classes distinguishers of $+\frac{q}{4}$ and $+\frac{q}{2}$ are complementary, as one gives away the Most Significant Bit (MSB) and the other the Least Significant Bit (LSB). However, we also notice that the bias $+\frac{3q}{4}$ can also be used alongside $+\frac{q}{2}$ or $+\frac{q}{4}$:

- The two-classes distinguisher of the bias $+\frac{3q}{4}$ gives the LSB, and thus can be paired with the bias $+\frac{q}{2}$ which gives the MSB.
- The three-classes distinguisher of the bias $+\frac{3q}{4}$ and the three-classes distinguisher of the bias $+\frac{q}{4}$ complement each other. The uncertainty in class $\{+1 \rightarrow 0\}$ for bias $+\frac{q}{4}$ is solved by the two classes $\{+0 \rightarrow 10\}$ and $\{+2 \rightarrow 00\}$ for bias $+\frac{3q}{4}$. Similarly, the uncertainty in the class $\{-1 \rightarrow 1\}$ linked to $+\frac{3q}{4}$ is solved by the two classes $\{+0 \rightarrow 01\}$ and $\{-2 \rightarrow 11\}$ linked to $+\frac{q}{4}$.

Frodo-640 can be breached by using ciphertext malleability with two Chosen Ciphertext queries and their associated oracle \mathcal{O}_{SCA} queries per couple of message bits. The length of the message in Frodo-640 is 128 bits. Full message recovery is possible in $128 + Q$ queries to \mathcal{O}_{SCA} , Q being the number of calls to \mathcal{O}_{SCA} to recover the Hamming Weight of the entire unbiased message. In terms of traces, for an implementation using 16-bit registers and parallelising the attacks, the complete attack costs 17 traces (only one trace is required to recover the Hamming Weight of the entire unbiased message). For an implementation using 8-bit registers, only 9 traces are required.

Remark 3.5. An attacker can decide to recover only half the message bits using any two-classes distinguisher in Frodo-640 and brute-force the remaining 64 bits of the message as all the two-classes distinguishers guarantee the recovery of at least one bit. This method allows for full message recovery with $64 + Q$ queries to the oracle (or 5 traces if 8-bit storage, 9 if 16-bit storage) and a brute-force complexity of 2^{64} in the worst case, which is feasible with current computer technology.

3.3.4. Frodo-976

The second security parameters set of FrodoKEM uses a different compression factor for message encoding with $d = 3$. Figure 1c gives a sector representation of the Decode function for this parameter. There are $2^d = 8$ possible outputs encoded in 3 bits. Table 5 highlights the impact of adding $\frac{q}{8}$, $\frac{q}{4}$, $\frac{q}{2}$, or $\frac{5q}{8}$ to the input of Decode on the Hamming Weight of its outputs.

Table 5: Impact of biasing the decoding function input with parameter $d = 3$ on the Hamming Weight of the output

Initial input	Initial mapping	$+\frac{q}{8}$	$+\frac{q}{4}$	$+\frac{q}{2}$	$+\frac{5q}{8}$
$[-\frac{q}{16}, \frac{q}{16}]$	000	+1	+1	+1	+2
$[\frac{q}{16}, \frac{3q}{16}]$	001	+0	+1	+1	+1
$[\frac{3q}{16}, \frac{5q}{16}]$	010	+1	+0	+1	+2
$[\frac{5q}{16}, \frac{7q}{16}]$	011	-1	+0	+1	-2
$[\frac{7q}{16}, \frac{9q}{16}]$	100	+1	+1	-1	+0
$[\frac{9q}{16}, \frac{11q}{16}]$	101	+0	+1	-1	-1
$[\frac{11q}{16}, \frac{13q}{16}]$	110	+1	-2	-1	+0
$[\frac{13q}{16}, \frac{15q}{16}]$	111	-3	-2	-1	-2

We have the following distinguishers for each bias used in Table 5:

- $+\frac{q}{8}$:
 - two-classes: $\{+1 \rightarrow \dots 0, (-3, -1, +0) \rightarrow \dots 1\}$
 - three-classes: $\{+1 \rightarrow \dots 0, +0 \rightarrow \dots 1, (-3, -1) \rightarrow \dots 1\}$
 - four-classes: $\{+1 \rightarrow \dots 0, +0 \rightarrow \dots 1, -3 \rightarrow 111, -1 \rightarrow 011\}$
- $+\frac{q}{4}$:
 - two-classes: $\{+1 \rightarrow \dots 0, (+0, -2) \rightarrow \dots 1\}$
 - three-classes: $\{+1 \rightarrow \dots 0, +0 \rightarrow \dots 1, -2 \rightarrow \dots 1\}$
- $+\frac{q}{2}$: two-classes: $\{-1 \rightarrow \dots 1, +1 \rightarrow \dots 0\}$
- $+\frac{5q}{8}$:
 - two-classes: $\{(-2, -1, +1) \rightarrow \dots 1, (+0, +2) \rightarrow \dots 0\}$
 - three-classes: $\{(-2, -1, +1) \rightarrow \dots 1, +2 \rightarrow \dots 0, +0 \rightarrow \dots 0\}, \{(-1, +1) \rightarrow \dots 1, -2 \rightarrow \dots 1, (+0, +2) \rightarrow \dots 0\}$
 - four-classes: $\{(-1, +1) \rightarrow \dots 1, -2 \rightarrow \dots 1, +2 \rightarrow \dots 0, +0 \rightarrow \dots 0\}, \{-1 \rightarrow \dots 1, +1 \rightarrow \dots 1, -2 \rightarrow \dots 1, (-2, +0) \rightarrow \dots 0\}$
 - five-classes: $\{-1 \rightarrow \dots 1, +1 \rightarrow \dots 1, +2 \rightarrow \dots 0, -2 \rightarrow \dots 1, +0 \rightarrow \dots 0\}$

As in Frodo-640, we cannot perform a successful attack in only one \mathcal{O}_{SCA} query per `Decode` as no column in Table 5 contains an eight-classes distinguisher. However, contrary to the attack on Frodo-640 where any couple of distinct biases is enough to recover the message bit values, it is not the case for Frodo-976.

For example, when considering biases $+\frac{q}{8}$ and $+\frac{q}{4}$, they do not complement each other to solve the uncertainty of the class $\{+1 \rightarrow ..0\}$ for the bias $+\frac{q}{8}$ or the uncertainty of the class $\{+1 \rightarrow .0.\}$ for the bias $+\frac{q}{4}$. However, a query with bias $+\frac{q}{2}$ gives the missing information. In fact, using the biases $+\frac{q}{8}$, $+\frac{q}{4}$ and $+\frac{q}{2}$ is equivalent to the Ravi et al. [28] attack as generalised in Section 3.2.

The targeted bit flip attack has a recovery efficiency of 1 bit per query. However, a better choice of biases allows for a more efficient bit recovery. We heuristically determined that querying \mathcal{O}_{SCA} for $+\frac{q}{4}$ and $+\frac{5q}{8}$ allows us to build an eight-classes distinguisher and consequently recover every bit of the message with only two oracle queries per 3 bits of message, as shown in Table 6, where the uncertainties left by the first query with bias $+\frac{q}{4}$ are solved by the query with bias $+\frac{5q}{8}$.

Table 6: Building a eight-classes distinguisher from $+\frac{q}{4}$ and $+\frac{5q}{8}$

$+\frac{q}{4}$	$+\frac{5q}{8}$	possibilities
$+1 \rightarrow .0.$	$-1 \rightarrow 101$	101
	$+1 \rightarrow 001$	001
	$+2 \rightarrow 0.0$	000
	$+0 \rightarrow 1.0$	100
$+0 \rightarrow 01.$	$-2 \rightarrow .11$	011
	$+2 \rightarrow 0.0$	010
$-2 \rightarrow 11.$	$-2 \rightarrow .11$	111
	$+0 \rightarrow 1.0$	110

A complete message recovery attack can be successfully performed against Frodo-976 with only $128 + Q$ queries to the \mathcal{O}_{SCA} oracle.

A peculiarity of Frodo-976 is that, since $d = 3$ does not divide 8 nor 16, the output of some calls to the `Decode` function will be split between two registers. For example, if 8-bit (resp. 16-bit) registers are used to store the message, in a 24-bit (resp. 48-bit) long word, that is, 3 registers, the first register is filled with the outputs of 3 (resp. 5) `Decode` functions. The same applies to the last register. However, the middle register is filled with the outputs of 4 (resp. 6) `Decode` functions that overlap with the first or last register. Although attacks on non-overlapping coefficients can be parallelised, the overlapping ones require specific investigation.

In the 8-bit register case, the overlap issue is represented in Figure 3, where the output of each `Decode` is stored in big endian. `Decode` 3 and `Decode` 6 affect two registers, with one in common (Register 2). `Decode` 1,4 and 7 can be attacked in parallel with two traces. The same applies to `Decode` 2,5 and 8. The number of traces required for `Decode` 3 and 6 can be reduced when specific biases are chosen. The biases $+\frac{q}{4}$ and $+\frac{3q}{4}$ do not affect the least significant bit of the output of `Decode`. As a consequence, choosing one of those two biases leads to a non-overlapping impact even if the `Decode` output is split between two registers. The number of traces required to attack the output of `Decode` 3 and 6 in a parallel manner is therefore 3. The total number of traces required to recover the message in an 8-bit implementation of Frodo-976 is 8 with our optimised bias choices.

Remark 3.6. The adaptation of the Ravi et al. [28] attack is also affected by the overlap issue. However, the bias $+\frac{q}{2}$ only affects the MSB and $+\frac{q}{4}$ the middle bit and the MSB. This implies that attacks against Decode 3 and Decode 6 can be parallelised for the two higher bits of Decode 3. Recovering the LSB of the output of Decode 3 cannot be parallelised with the recovery of the last bit of Decode 6. Thus, it requires an extra trace, for a total of 10.

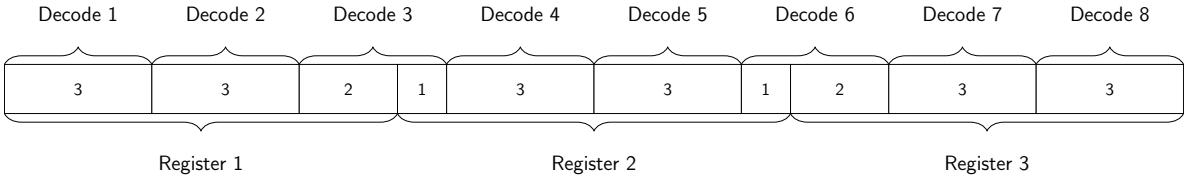


Figure 3: Overlap issue with 8 decodings and their 24-bit long output split on 3 registers

In the 16-bit register case, the overlap issue is represented in Figure 4. Decode 6 and Decode 11 affect two registers and thus would require 4 traces. However, the middle register has only four complete Decode compared to the five of the first and last registers. Using the observation that the biases $+\frac{q}{4}$ and $+\frac{3q}{4}$ do not affect the least significant bit of the output of Decode, we reduce the number of required traces. We parallelise as follows: Decode 1, 7 and 12 (two traces), Decode 2, 8 and 13 (two traces), Decode 3, 9 and 14 (two traces), Decode 4, 10 and 15 (two traces). We perform the two attacks on Decode 5 and 16 in parallel with only one attack against Decode 11 using the aforementioned biases (two traces). With the three remaining traces to sort the overlapping Decode instances, and the original trace, the complexity of the parallelised attack against Frodo-976 implemented with 16-bit registers is 14 traces.

Remark 3.7. Remark 3.6 also applies to the 16-bit register case for the adaptation of Ravi et al. [28] attack. The parallelisation in this case requires not one but two extra traces, for a total of 19.

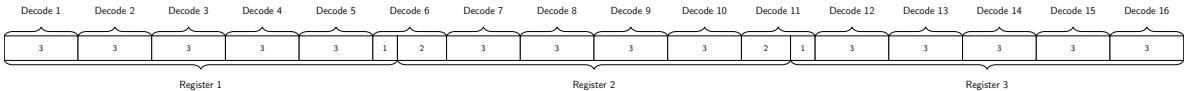


Figure 4: Overlap issue with 16 decodings and their 48-bit long output split on 3 registers

3.3.5. Frodo-1344

The highest security level of FrodoKEM uses $d = 4$. With this parameter, the Decode function maps to $2^d = 16$ possible outputs encoded on 4 bits. Table 7 shows the impact of adding $+\frac{2^j}{16}$ with $j \in \llbracket 0, 3 \rrbracket$ or $+\frac{5q}{8}$ to the input of Decode.

Table 7: Impact of biasing the decoding function input with parameter $d = 4$ on the Hamming Weight of the output

Initial input	Initial mapping	$+\frac{q}{16}$	$+\frac{q}{8}$	$+\frac{q}{4}$	$+\frac{q}{2}$	$+\frac{5q}{8}$
$[-\frac{q}{32}, \frac{q}{32}]$	0000	+1	+1	+1	+1	+2
$[\frac{q}{32}, \frac{3q}{32}]$	0001	+0	+1	+1	+1	+2
$[\frac{3q}{32}, \frac{5q}{32}]$	0010	+1	+0	+1	+1	+1
$[\frac{5q}{32}, \frac{7q}{32}]$	0011	-1	+0	+1	+1	+1
$[\frac{7q}{32}, \frac{9q}{32}]$	0100	+1	+1	+0	+1	+2
$[\frac{9q}{32}, \frac{11q}{32}]$	0101	+0	+1	+0	+1	+2
$[\frac{11q}{32}, \frac{13q}{32}]$	0110	+1	-1	+0	+1	-2
$[\frac{13q}{32}, \frac{15q}{32}]$	0111	-2	-1	+0	+1	-2
$[\frac{15q}{32}, \frac{17q}{32}]$	1000	+1	+1	+1	-1	+0
$[\frac{17q}{32}, \frac{19q}{32}]$	1001	+0	+1	+1	-1	+0
$[\frac{19q}{32}, \frac{21q}{32}]$	1010	+1	+0	+1	-1	-1
$[\frac{21q}{32}, \frac{23q}{32}]$	1011	-1	+0	+1	-1	-1
$[\frac{23q}{32}, \frac{25q}{32}]$	1100	+1	+1	-2	-1	+0
$[\frac{25q}{32}, \frac{27q}{32}]$	1101	+0	+1	-2	-1	+0
$[\frac{27q}{32}, \frac{29q}{32}]$	1110	+1	-3	-2	-1	-2
$[\frac{29q}{32}, \frac{31q}{32}]$	1111	-4	-3	-2	-1	-2

The biases $+\frac{q}{16}$, $+\frac{q}{8}$, $+\frac{q}{4}$ and $+\frac{q}{2}$ are used to perform the Ravi et al. [28] targeted bit flip attack. The efficiency of this attack remains at 1 bit per chosen ciphertext query. We heuristically determined that the minimum number of biases for a complete recovery of the output of $\text{Decode}(\alpha, d)$ is 3. Combining the biases $+\frac{q}{16}$, $+\frac{q}{4}$ and $+\frac{5q}{8}$ allows us to build a sixteen-classes distinguisher and perform the attack successfully with only three Chosen Ciphertext queries to \mathcal{O}_{SCA} per Decode . Table 8 describes the resolution of the uncertainties when the first bias is $+\frac{q}{4}$, the second is $+\frac{5q}{8}$ and the last² $+\frac{q}{16}$.

There are 95 possible choices of 3 biases to perform a complete Decode output recovery. Frodo-1344 can be attacked for complete message recovery in only $192 + Q$ Chosen Ciphertext queries with the help of the \mathcal{O}_{SCA} oracle. For an implementation using 16-bit registers, thanks to the parallelisation of the attacks, only 13 traces are required for a complete message recovery. For an implementation using 8-bit registers, only 7 traces are required.

3.4. Observations for higher values of d

In this work, we determined heuristically the values of the optimal biases. However, we observed patterns in the choice of these biases. In this section, we discuss some of them and how they could affect message recovery for higher values of the parameter d . Other observations are conjectures; we leave a more thorough study for future work.

²The order of the biases is for readability and has no impact on the attack

Table 8: Building a sixteen-classes distinguisher from $+\frac{q}{4}$, $+\frac{5q}{8}$ and $+\frac{q}{16}$

$+\frac{q}{4}$	$+\frac{5q}{8}$	$+\frac{q}{16}$	possibilities
$+1 \rightarrow .0\dots$	$+2 \rightarrow 0\ .0.$	$+1 \rightarrow \dots 0$	0000
		$+0 \rightarrow \dots 01$	0001
	$+1 \rightarrow 001.$	$+1 \rightarrow \dots 0$	0010
		$-1 \rightarrow .011$	0011
		$+0 \rightarrow 1\ .0.$	1000
		$+0 \rightarrow \dots 01$	1001
	$-1 \rightarrow 101.$	$+1 \rightarrow \dots 0$	1010
		$-1 \rightarrow .011$	1011
$+0 \rightarrow 01\dots$	$+2 \rightarrow 0\ .0.$	$+1 \rightarrow \dots 0$	0100
		$+0 \rightarrow \dots 01$	0101
	$-2 \rightarrow .11.$	$+1 \rightarrow \dots 0$	0110
		$-2 \rightarrow 0111$	0111
$-2 \rightarrow 11\dots$	$+0 \rightarrow 1\ .0.$	$+1 \rightarrow \dots 0$	1100
		$+0 \rightarrow \dots 01$	1101
	$-2 \rightarrow .11.$	$+1 \rightarrow \dots 0$	1110
		$-4 \rightarrow 1111$	1111

Theorem 3.8. Any choice of biases for complete message recovery includes at least one bias $+\frac{qk}{2^d}$ with k an odd number.

Proof:

We define the binary decomposition of a bias k as $k = \sum_{i=0}^{d-1} 2^i k_i$. Similarly, the binary decomposition of the output of $\text{Decode}(\alpha, d)$ is $m = \sum_{i=0}^{d-1} 2^i m_i$.

$$\text{Decode}(\alpha + \frac{qk}{2^d}, d) \equiv \text{Decode}(\alpha, d) + k \equiv \text{Decode}(\alpha, d) + k_0 + 2k_1 + \dots + 2^{d-1}k_{d-1} \bmod 2^d. \quad (7)$$

With the decomposition described in Equation 7, the calculation of the impact of the bias on the output of Decode can be performed by taking the impact of adding each $2^i k_i$ one after the other. Consequently, we can take each k_i separately, compute a temporary Δ according to Equation 5, do the same with k_{i+1} but taking into account the previous modification performed if $k_i = 1$ to compute the next temporary Δ . Given the representation in sectors, this is equivalent to performing a first rotation of $2^i k_i$ sectors, computing the temporary Δ , and then performing the next rotation by $2^{i+1} k_{i+1}$ to add the results.

Let us assume that $k_0 = 0$ in Equation 7, i.e. k is an even number. Then, the impact of the bias on the output bit m_0 is null as $\text{Decode}(\alpha, d) = \text{Decode}(\alpha, d) + 0$. Regardless of the value of the next k_i

such as $i > 0$, they do not affect the output bit m_0 . According to Equation 5, the next $2^i k_i$ can affect only m_1 to m_{d-1} . As a consequence, there is always an uncertainty in the value of m_0 as it is never altered by the bias added to the input of `Decode`, and therefore has never an impact on the variation of the Hamming Weight of m . To build a 2^d -classes distinguisher, it is mandatory to have at least one bias $+\frac{q^k}{2^d}$ with $k_0 = 1$. \square

A comparison between the efficiency of the Ravi et al. [28] targeted bit flip attack and a more optimal choice of biases is given in Figure 5. We heuristically determined the minimal number of biases required for a complete recovery by performing an exhaustive search for increasing values of the number of biases.

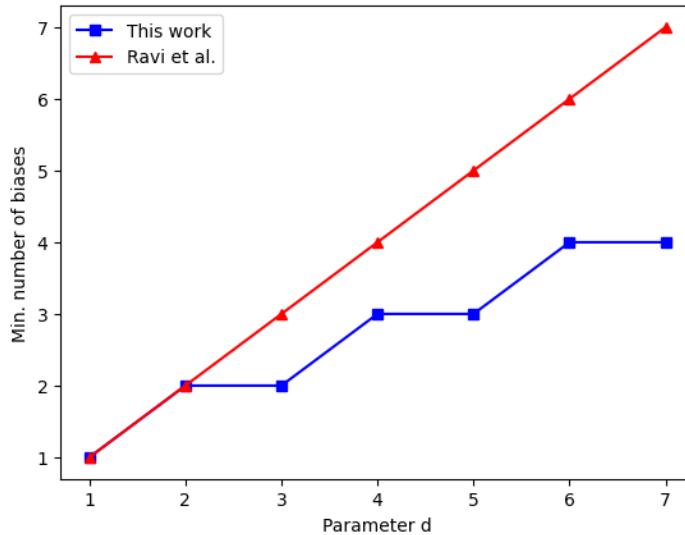


Figure 5: Comparison of the number of biases required for a complete output recovery between Ravi et al. [28] and this work

4. Countermeasure

4.1. Overall strategy

Ravi et al. [28] claim that generic countermeasures such as shuffling or masking cannot protect against their attack. In this section, we propose reusing the attack as a defence mechanism. The rationale for the design is to apply the attack in a random but controlled manner. This will randomly flip the message bits to a point where an attacker cannot differentiate between his own attack and the countermeasure, thus counteracting the attack locally.

This countermeasure has an effect on the Hamming Weight of the message similar to the generic masking countermeasure. However, contrary to masking where each share is processed separately, in our case only one “share” is processed. This can reduce the cost of the implementation, as masking is known to have significant overhead due to its multiple shares to process.

In this section, we present the countermeasure for the LPR framework and also detail its application to specific algorithms.

4.2. Application to the decryption

To perform the countermeasure, we randomly choose a bias among a predefined set of biases and add it to the input of the `Decode` function. This can be performed by sampling random bits and using a map between those bits and the set of biases.

We can reuse existing parts of the LPR framework to perform the two operations (sampling and mapping). Random sampling is already performed in LPR during the message sampling in the `KEM.Encaps` procedure as described in Algorithm 2. For the mapping, it is performed by the `Encode` function used in the `PKE.Encrypt` procedure. The message is then added to the received ciphertext \mathbf{v} in the same way as the message is added to the temporary ciphertext of \mathbf{v} during the `PKE.Encrypt` procedure.

In Algorithm 3 a colour code is used to highlight the provenances of each part of the countermeasure. In **blue** are the functions reused from `KEM.Encaps`, in **green** are the ones reused from `PKE.Encrypt`.

Algorithm 3: Decryption with countermeasure

```

1 Procedure PKE.DecryptWithCountermeasure( $ct, sk$ ):
2    $mask \leftarrow \mathcal{U}(\mathcal{B}^{32})$ 
3    $\mathbf{vm} = \mathbf{v} + \text{Encode}(mask, d)$ 
4    $\mathbf{xm}' = (\mathbf{vm} - \mathbf{u} \times \mathbf{s})$ 
5    $mm = \text{Decode}(\mathbf{xm}', d)$ 
6   return  $mm, mask$ 

```

Due to ciphertext malleability, we have the following theorem:

Theorem 4.1. Let m be the message output of `PKE.Decrypt`(ct, sk). Let mm be the message output and $mask$ the mask output of `PKE.DecryptWithCountermeasure`(ct, sk). Then, for each sub-message m_i , mm_i and $mask_i$ of length d bits,

$$mm_i = mask_i + m_i \bmod 2^d. \quad (8)$$

Proof:

This is a direct consequence of ciphertext malleability. Let us consider the `Encode` and `Decode` functions for the parameter d . We use the sector representation for `Decode`. The function `Encode` maps the bit words of value k between 0 and $2^d - 1$ to the corresponding values $\frac{qk}{2^d}$. When used as biases on the input of `Decode`, these values rotate the sectors, as discussed in Section 3.2. As each sector covers $\frac{qk}{2^d}$ values, applying any of the aforementioned biases will rotate the sectors in a way that retains the set of the bounds of every sector. However, each sector now corresponds to its previous value $+k \bmod 2^d$. We then take $k = mask_i$ and reapplied this proof to the other sub-messages of $mask$. \square

Remark 4.2. If the parameter d used for `Encode` and `Decode` is 1, then $mm = m \oplus \text{mask}$ as $\forall (a, b) \in \{0, 1\}, a + b \bmod 2 \Leftrightarrow a \oplus b$. This is the case for ML-KEM and SABER.

4.3. Extension to the remaining decapsulation

Due to the Fujisaki-Okamoto Transform, the message is later re-encrypted during the `KEM.Decaps` procedure. In this section, we discuss how to extend our countermeasure to the remainder of the decapsulation.

4.3.1. Seed generation

The message is used to generate the random seed used in `PKE.Encrypt`. This is done by using a PRF on the message and a token derived from the public key. In both FrodoKEM [22], ML-KEM [23], and SABER [4], this is performed with a call to a Keccak instance taking the message concatenated with a hash of the public key as input. This is central within the correctness of the FOT. If the message is altered, the seed generated during the decapsulation will be significantly different from the one used during the encapsulation, resulting in a different ciphertext altogether. Therefore, we need to ensure the correctness of the seed.

The application of our countermeasure results in a masked message. This masking is an arithmetic masking modulo 2^d , or a boolean masking if $d = 1$. Seed generation uses a hash function, implying computations with boolean logic. To perform it in a secure manner, if $d \neq 1$, we must first convert the masking from arithmetic to boolean (A2B conversion). Once done, this can be fed into a secure masked hash function. However, such an implementation of the hash function is costly but necessary.

4.3.2. Re-encryption

To further extend the countermeasure, we propose to perform re-encryption on the masked message but with the correct seed. The calculations of \mathbf{u} and \mathbf{v}' during `PKE.Encrypt` are therefore not altered by our countermeasure. However, we have $\mathbf{v}'' = \mathbf{v}' + \text{Encode}(mm)$, resulting in a different ciphertext.

After the re-encryption is performed, the generated ciphertext is compared with the received one. To ensure correctness, we have to “correct” \mathbf{v}'' so it corresponds to \mathbf{v} . This correction depends on the linearity of the `Encode` function that is tied to q . It is therefore primitive-dependent. If 2^d divides q , then the rounding in Equation 3 is not necessary, since both β and $\frac{q}{2^d}$ are integers. Multiplying an integer by a constant scalar is linear.

Both SABER [4] and FrodoKEM [22] have this property. In both cases, the modulo used within `Encode` is a multiple of the compression factor, and consequently the rounding is not needed. This is not the case for ML-KEM [23].

Linear encoding As the function `Encode` is linear, to compute the correct ciphertext one can use the encoding of the mask on the generated ciphertext. For FrodoKEM [22], when performing the re-encryption during the decapsulation, instead of lines 11 and 12 in Algorithm 1, we compute³

$$\mathbf{v} = \mathbf{v}' + \text{Encode}(mm, d) - \text{Encode}(\text{mask}, d). \quad (9)$$

³With FrodoKEM [22] notation: $\mathbf{C}' = \mathbf{V} + \text{Frodo.Encode}(mm) - \text{Frodo.Encode}(\text{mask})$.

For SABER [4], the `Encode` function is not defined but its effect is still used. An interesting detail is that the encoding is not added but subtracted. Thus, when performing the re-encryption during the decapsulation, instead of lines 11 and 12 in Algorithm 1, we compute⁴

$$\mathbf{v} = \mathbf{v}' - \text{Encode}(mm, d) + \text{Encode}(mask, d). \quad (10)$$

Non-linear encoding ML-KEM [23] uses a Solinas/Crandall prime. Therefore, it cannot be divided by 2 and as a result, the encoding is not linear. There are several options. We take inspiration from existing work on masked implementations of ML-KEM [9, 14].

The first idea is to ensure correctness by considering the encoding to be linear and then to further correct in specific cases. When performing the re-encryption during the decapsulation, instead of lines 11 and 12 in Algorithm 1 we compute⁵

$$\mathbf{v} = \mathbf{v}' + \text{Encode}(mm, 1) + \text{Encode}(mask, 1). \quad (11)$$

However, as shown in Table 9, there is a small bias when both *mask* and *mm* are equal to 1. Table 9 uses the following legend:

- white : The result is equal with $\text{Encode}(m = mask \oplus mm, 1)$
- white : The result differs from $\text{Encode}(m, 1)$

Table 9: Comparing $\text{Encode}(mm, 1) + \text{Encode}(mask, 1)$ with $\text{Encode}(m, 1)$

	$mm_i = 0$	$mm_i = 1$
$mask_i = 0$	0+0	1665+0
$mask_i = 1$	0+1665	1665+1665=1 \neq 0

There are two ways to further correct this bias. The first relies on the masked implementation of Heinz et al. [14]. They propose to subtract the logic AND of the mask and the masked message from the computations of \mathbf{v}' before performing the ciphertext comparison. However, if not performed securely, this AND will reveal information on the message.

The second way is inspired by the masked implementation of Bos et al. [9]. Instead of comparing the ciphertexts by compressing the generated ciphertext and checking for a strict equality between ct and ct' , they propose to compare the generated ciphertext directly with the decompression of the received one. However, since compression and decompression are lossy, they do not check for a strict equality but rather that the maximum distance between both ciphertexts does not exceed a specific threshold. To compensate for the bias, we suggest increasing the threshold for ciphertext \mathbf{v} by one.

⁴With SABER [4] notation: $c_m = (v' + h_1 - 2^{\epsilon_p-1}mm + 2^{\epsilon_p-1}mask \bmod p) \gg (\epsilon_p - \epsilon_T)$.

⁵With ML-KEM [23] notation: $v' = \vec{t}^\perp \cdot \vec{r} + e_1 + \text{Decompress}_q(mm, 1) + \text{Decompress}_q(mask, 1)$.

We propose a new alternative to these methods. Instead of adding the decompression of the mask and then correct the resulting bias, we propose the following formula⁶:

$$\mathbf{v} = \mathbf{v}' + \text{Encode}(mm, 1) + (-1)^{mm} \text{Encode}(mask, 1). \quad (12)$$

This results in Table 10, with the same legend as Table 9.

Table 10: Comparison between $\text{Encode}(mm, 1) + (-1)^{mm} \text{Encode}(mask, 1)$ and $\text{Encode}(m, 1)$

	$mm_i = 0$	$mm_i = 1$
$mask_i = 0$	0+0	1665-0
$mask_i = 1$	0+1665	1665-1665=0

4.4. Scalability

Scalability depends on the linearity of the encoding used. For FrodoKEM [22] and SABER [4], our countermeasure can be easily scaled to simulate higher orders of masking by simply generating new masks, encoding them, and adding them to the ciphertext. As a result, we end up with a higher order of masking for the message while still computing “one share”.

For ML-KEM [23], the scalability is less trivial as the encoding is non-linear. Adding $\lceil \frac{q}{2} \rceil$ several times in a row creates a growing bias as $2\lceil \frac{q}{2} \rceil \neq q$ due to the parity of q . The probability of having a failed bitflip resulting in an incorrect masking is low, but grows with this bias. One way of slowing the growth of this bias is to alternate between adding and subtracting $\lceil \frac{q}{2} \rceil$. This does not alleviate the growing complexity of the adaptations of Equations 11 and 12.

4.5. Implementation recommendations

4.5.1. Mask generation

To generate $mask$, we use the same constraints as for the generic masking countermeasure. The share must be generated with a cryptographically secure RNG or a TRNG. In the *pqm4* implementation [18], a hardware RNG is used to generate the message, we recommend using the same to generate $mask$.

4.5.2. On-the-fly encoding

A major drawback of post-quantum primitives is the size of the computations. Hence, we recommend implementing the function `Encode` (used at Line 3 of Algorithm 3) following an “On-the-fly” philosophy, i.e. encoding a coefficient of $mask$ and adding it to \mathbf{v} at the proper place before encoding the next coefficient. A similar approach should be preferred for other calls to the function `Encode`, e.g. during the re-encapsulation for the encoding of mm . This design is inspired by the work of Bos et al. [8], where they use such an approach to optimise the memory usage of a FrodoKEM implementation on

⁶With ML-KEM [23] notation: $v' = \vec{t}^\perp \cdot \vec{r} + e_1 + \text{Decompress}_q(mm, 1) + (-1)^{mm} \text{Decompress}_q(mask, 1)$.

Cortex M4. As a consequence, we avoid storing encodings in memory, thus reducing memory stack consumption. This also benefits side-channel resistance, as attacking internal registers is much harder than attacking memory storage. The drawback is an increase in the complexity of implementation in time. In FrodoKEM [22], ML-KEM [23] and SABER [4], the function `Encode`, when applied to the message and thus to mm or $mask$, is a simple scalar multiplication.

4.6. Vulnerability and compatibility with other countermeasures

Although our countermeasure acts as a natural countermeasure to the use of ciphertext malleability [28], it is not a generic standalone countermeasure, as it does not affect the secret. Another issue resulting from the reuse of the primitive functions is that existing side-channel attacks against these functions, especially against the function `Encode` [2, 33], can circumvent the countermeasure. Thus, for a fully secured implementation, our method must be combined with other countermeasures.

As our countermeasure can be performed with minimal new code and, from a high-level perspective, reuses existing functions within LPR-based primitives, it is fully compatible with other secured implementations such as masking and/or shuffling. For a masked implementation, our countermeasure can be used to artificially increase the masking order of the message in the implementation.

We recommend combining our countermeasure with the generic shuffling countermeasure [15, 35], especially for calls to function `Decode` and function `Encode`. Shuffling is considered an effective countermeasure to protect the function `Encode` [29, 2].

In a way, our countermeasure forces the attacker that was targeting the decoding function, for which generic countermeasures were proved ineffective, to target the encoding function where a generic countermeasure, shuffling, is known to be effective.

Remark 4.3. Most of the state-of-the-art masked implementations of ML-KEM perform `Decode` by first converting from arithmetic to boolean logic and then using boolean logic to compute the decoding. However, there are alternatives which compute decoding using arithmetic and then convert the result into boolean logic, such as [6]. Our countermeasure can be used in this case to perform a cheaper arithmetic to boolean conversion by simply unmasking rather than converting. As we provide a first-order boolean masking with our countermeasure, unmasking in arithmetic does not expose the secret, but only one boolean share of it.

4.7. Discussion on the impact on other attacks

Some attacks aim to recover the entire message using altered ciphertexts such as the one studied in Section 3. From a generic point of view, any attack that requires complete control over the ciphertext v and using the same oracle \mathcal{O}_{SCA} described in our paper is impacted.

This applies to several attacks against ML-KEM categorised in [29]:

- *Binary Plaintext-Cheking* attacks [30, 32, 26, 5] use sparse ciphertexts to target single coefficients of the secret,
- *Parallel Plaintext-Checking* attacks [27, 34] use sparse ciphertexts as well but target multiple coefficients of the secret at the same time.

5. Conclusion

In this paper we recall an attack from Ravi et al. [28] aimed at recovering the message in LPR-based cryptography primitives. We extended their work by detailing the attack on a post-quantum KEM, namely FrodoKEM, in Section 3. We also generalised their initial attack to any value of a parameter d and discussed more optimal choices to reduce the number of traces required for a complete message recovery. We presented a novel countermeasure to this specific attack in Section 4. The countermeasure is based on the same ciphertext malleability principle as the attack. We discussed the extension of the countermeasure to the entirety of the KEM.Decaps procedure, as well as its compatibility with other generic countermeasures and its scalability for several post-quantum primitives. As the countermeasure efficiency is local, we recommended its deployment in conjunction with a shuffling countermeasure.

The Ravi et al. [28] attack principle is also used in several *Deep Learning-SCA* (DL-SCA). It will be interesting to see how our countermeasure performs against a DL-SCA opponent. However, this will be implementation specific, as DL-SCA attacks tend to target specific *determiner-leakages* due to the specific methods of computation used in some implementations.

Acknowledgments The author thanks Yoan Rougeolle for his help with the experimental part of this work.

This work was realised thanks to the grant 2022156 from the Appel à projets 2022 thèses AID Cifre-Défense by the Agence de l’Innovation de Défense (AID), Ministère des Armées (French Ministry of Defence).

References

- [1] Alagic G, Apon D, Cooper D, Dang Q, Dang T, Kelsey J, Lichtinger J, Miller C, Moody D, Peralta R, et al. Status report on the third round of the NIST post-quantum cryptography standardization process. *US Department of Commerce, NIST*, 2022. doi:10.6028/NIST.IR.8413-upd1.
- [2] Amiet D, Curiger A, Leuenberger L, Zbinden P. Defeating NewHope with a Single Trace. In: Ding J, Tillich JP (eds.), *Post-Quantum Cryptography*. Springer International Publishing, Cham. ISBN 978-3-030-44223-1, 2020 pp. 189–205.
- [3] Avanzi R, Bos J, Ducas L, Kiltz E, Lepoint T, Lyubashevsky V, Schanck JM, Schwabe P, Seiler G, Stehlé D. CRYSTALS-Kyber Algorithm Specifications And Supporting Documentation. *pq-crystals*, 2021.
- [4] Basso A, Mera JMB, D’Anvers JP, Karmakar A, Roy SS, Beirendonck MV, Vercauteren F. SABER: Mod-LWR based KEM (round 3 submission). *Submission to the NIST’s postquantum cryptography standardization process*, 2020.
- [5] Băetu C, Durak FB, Huguenin-Dumittan L, Talayhan A, Vaudenay S. Misuse attacks on post-quantum cryptosystems. In: *Advances in Cryptology–EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, Darmstadt, Germany, May 19–23, 2019, Proceedings, Part II 38. Springer, 2019 pp. 747–776.

- [6] Berthet PA, Rougeolle Y, Tavernier C, Danger JL, Sauvage L. Securing Lattice-Based KEMs with Code-Based Masking: A Theoretical Approach. *Cryptology ePrint Archive*, Paper 2023/1220, 2023. <https://eprint.iacr.org/2023/1220>, URL <https://eprint.iacr.org/2023/1220>.
- [7] Bos JW, Friedberger S, Martinoli M, Oswald E, Stam M. Fly, you fool! Faster Frodo for the ARM Cortex-M4. *Cryptology ePrint Archive*, Report 2018/1116, 2018. URL <https://eprint.iacr.org/2018/1116>.
- [8] Bos JW, Bronchain O, Custers F, Renes J, Verbakel D, van Vredendaal C. Enabling FrodoKEM on Embedded Devices. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023. **2023**(3):74–96. doi:10.46586/tches.v2023.i3.74-96. URL <https://tches.iacr.org/index.php/TCIES/article/view/10957>.
- [9] Bos JW, Gourjon M, Renes J, Schneider T, Van Vredendaal C. Masking kyber: First-and higher-order implementations. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2021. pp. 173–214. doi:10.46586/tches.v2021.i4.173-214.
- [10] Canto AC, Kaur J, Kermani MM, Azarderakhsh R. Algorithmic security is insufficient: A comprehensive survey on implementation attacks haunting post-quantum security. *arXiv preprint arXiv:2305.13544*, 2023.
- [11] Chari S, Jutla CS, Rao JR, Rohatgi P. Towards sound approaches to counteract power-analysis attacks. In: *Advances in Cryptology—CRYPTO’99: 19th Annual International Cryptology Conference* Santa Barbara, California, USA, August 15–19, 1999 Proceedings 19. Springer, 1999 pp. 398–412. doi: 10.1007/3-540-48405-1_26.
- [12] Chen L, Chen L, Jordan S, Liu YK, Moody D, Peralta R, Perlner RA, Smith-Tone D. Report on post-quantum cryptography, volume 12. US Department of Commerce, National Institute of Standards and Technology ..., 2016.
- [13] Fujisaki E, Okamoto T. Secure integration of asymmetric and symmetric encryption schemes. In: *Annual international cryptology conference*. Springer, 1999 pp. 537–554. doi:10.1007/3-540-48405-1_34.
- [14] Heinz D, Kannwischer MJ, Land G, Pöppelmann T, Schwabe P, Sprenkels A. First-Order Masked Kyber on ARM Cortex-M4. *Cryptology ePrint Archive*, Paper 2022/058, 2022. <https://eprint.iacr.org/2022/058>, URL <https://eprint.iacr.org/2022/058>.
- [15] Herbst C, Oswald E, Mangard S. An AES smart card implementation resistant to power analysis attacks. In: *International conference on applied cryptography and network security*. Springer, 2006 pp. 239–252. doi:10.1007/11767480_16.
- [16] Howe J, Oder T, Krausz M, Güneysu T. Standard Lattice-Based Key Encapsulation on Embedded Devices. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018. **2018**(3):372–393. doi: 10.13154/tches.v2018.i3.372-393. URL <https://tches.iacr.org/index.php/TCIES/article/view/7279>.
- [17] Ishai Y, Sahai A, Wagner D. Private circuits: Securing hardware against probing attacks. In: *Advances in Cryptology-CRYPTO 2003: 23rd Annual International Cryptology Conference*, Santa Barbara, California, USA, August 17-21, 2003. Proceedings 23. Springer, 2003 pp. 463–481. doi:10.1007/978-3-540-45146-4_27.
- [18] Kannwischer MJ, Petri R, Rijneveld J, Schwabe P, Stoffelen K. PQM4: Post-quantum crypto library for the ARM Cortex-M4. <https://github.com/mupq/pqm4/tree/Round3>.

- [19] Kocher PC. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Advances in Cryptology—CRYPTO’96: 16th Annual International Cryptology Conference Santa Barbara, California, USA August 18–22, 1996 Proceedings 16. Springer, 1996 pp. 104–113. doi:10.1007/3-540-68697-5_9.
- [20] Lyubashevsky V, Peikert C, Regev O. On Ideal Lattices and Learning with Errors over Rings. In: Gilbert H (ed.), Advances in Cryptology – EUROCRYPT 2010. Springer Berlin Heidelberg, Berlin, Heidelberg. ISBN 978-3-642-13190-5, 2010 pp. 1–23.
- [21] Mamiya H, Miyaji A, Morimoto H. Efficient countermeasures against RPA, DPA, and SPA. In: International workshop on cryptographic hardware and embedded systems. Springer, 2004 pp. 343–356. doi:10.1007/978-3-540-28632-5_25.
- [22] Naehrig M, Alkim E, Bos J, Ducas L, Easterbrook K, LaMacchia B, Longa P, Mironov I, Nikolaenko V, Peikert C, et al. FrodoKEM. *National Institute of Standards and Technology* (2017), 2017.
- [23] NIST. Module-Lattice-Based Key-Encapsulation Mechanism Standard. *NIST FIPS*, 2023. doi:10.6028/NIST.FIPS.203.ipd.
- [24] Plos T, Hutter M, Herbst C. Enhancing side-channel analysis with low-cost shielding techniques. In: Proceedings of Austrochip. 2008 pp. 90–95. URL <http://www.mhutter.org/papers/Plos2008EnhancingSideChannelAnalysis.pdf>.
- [25] Prouff E, Rivain M. Masking against side-channel attacks: A formal security proof. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. Springer, 2013 pp. 142–159. doi:10.1007/978-3-642-38348-9_9.
- [26] Qin Y, Cheng C, Zhang X, Pan Y, Hu L, Ding J. A systematic approach and analysis of key mismatch attacks on lattice-based NIST candidate KEMs. In: Advances in Cryptology—ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part IV 27. Springer, 2021 pp. 92–121.
- [27] Rajendran G, Ravi P, D’anvers JP, Bhasin S, Chattopadhyay A. 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*, 2023.
- [28] Ravi P, Bhasin S, Roy SS, Chattopadhyay A. On exploiting message leakage in (few) NIST PQC candidates for practical message recovery attacks. *IEEE Transactions on Information Forensics and Security*, 2021. **17**:684–699. doi:10.1109/TIFS.2021.3139268.
- [29] Ravi P, Chattopadhyay A, D’Anvers JP, Baksi A. Side-channel and fault-injection attacks over lattice-based post-quantum schemes (Kyber, Dilithium): Survey and new results. *ACM Transactions on Embedded Computing Systems*, 2024. **23**(2):1–54.
- [30] Ravi P, Roy SS, Chattopadhyay A, Bhasin S. Generic side-channel attacks on CCA-secure lattice-based PKE and KEMs. *IACR transactions on cryptographic hardware and embedded systems*, 2020. pp. 307–335.
- [31] Saarinen MJO. Arithmetic coding and blinding countermeasures for lattice signatures: Engineering a side-channel resistant post-quantum signature scheme with compact signatures. *Journal of Cryptographic Engineering*, 2018. **8**(1):71–84. doi:10.1007/s13389-017-0149-6.
- [32] Shen M, Cheng C, Zhang X, Guo Q, Jiang T. Find the bad apples: An efficient method for perfect key recovery under imperfect sca oracles—a case study of kyber. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023. pp. 89–112.

- [33] Sim BY, Kwon J, Lee J, Kim IJ, Lee TH, Han J, Yoon H, Cho J, Han DG. Single-Trace Attacks on Message Encoding in Lattice-Based KEMs. *IEEE Access*, 2020. **8**:183175–183191. doi:10.1109/ACCESS.2020.3029521.
- [34] Tanaka Y, Ueno R, Xagawa K, Ito A, Takahashi J, Homma N. Multiple-valued plaintext-checking side-channel attacks on post-quantum kems. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023. **2023**(3):473–503.
- [35] Veyrat-Charvillon N, Medwed M, Kerckhof S, Standaert FX. Shuffling against side-channel attacks: A comprehensive study with cautionary note. In: Advances in Cryptology–ASIACRYPT 2012: 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings 18. Springer, 2012 pp. 740–757. doi:10.1007/978-3-642-34961-4_44.