



Assessment of the Key-Reuse Resilience of NewHope

Aurélie Bauer¹, Henri Gilbert^{1,2}, Guénaél Renault^{1,3(✉)}, and Mélissa Rossi^{4,5}

¹ ANSSI, Paris, France

{aurelie.bauer,henri.gilbert,guenael.renault}@ssi.gouv.fr

² UVSQ, Versailles, France

³ Sorbonne Université, CNRS, Inria, Laboratoire d'Informatique de Paris 6, LIP6, Équipe PolSys, Paris, France

⁴ Thales, Gennevilliers, France

⁵ Département d'informatique de l'Ecole normale supérieure, CNRS, PSL Research University, Inria, Paris, France
melissa.rossi@ens.fr

Abstract. *NewHope* is a suite of two efficient Ring-Learning-With-Error based key encapsulation mechanisms (KEMs) that has been proposed to the NIST call for proposals for post-quantum standardization. In this paper, we study the security of *NewHope* when an active adversary takes part in a key establishment protocol and is given access to an oracle, called *key mismatch oracle*, which indicates whether her guess of the shared key value derived by the party targeted by the attack is correct or not. This attack model turns out to be relevant in private key reuse situations since an attacker may then be able to access such an oracle repeatedly – either directly or using faults or side channels, depending on the considered instance of *NewHope*. Following this model we show that, by using *NewHope* recommended parameters, several thousands of queries are sufficient to recover the full private key with high probability. This result has been experimentally confirmed using *Magma* CAS implementation. While the presented *key mismatch oracle* attacks do not break any of the designers' security claims for the *NewHope* KEMs, they provide better insight into the resilience of these KEMs against key reuse. In the case of the CPA-KEM instance of *NewHope*, they confirm that key reuse (e.g. key caching at server side) should be strictly avoided, even for an extremely short duration. In the case of the CCA-KEM instance of *NewHope*, they allow to point out critical steps inside the CCA transform that should be carefully protected against faults or side channels in case of potential key reuse.

Keywords: PQ-crypto · Lattice based cryptography · Active attack · Side channels

This research has been partially funded by ANRT under the program CIFRE 2016/1583. We acknowledge the support of the French Programme d'Investissement d'Avenir under national project RISQ P141580. This work is also partially supported by the European Union PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701).

1 Introduction

The insecurity of the main asymmetric cryptosystems (RSA, (EC)DLP) in front of a potential quantum computer has led the cryptographic community to investigate new quantum resistant primitives. In 2016, NIST has initiated a process to develop and standardize one or more public-key cryptographic algorithms which are supposed to be quantum safe. Cryptosystems based on lattices represent one of the most promising directions for such systems.

Key Encapsulation Mechanisms (or KEMs) are one of the most important asymmetric cryptographic primitives. The NIST call specifically asks for quantum resistant KEM proposals in order to replace number theory based Diffie-Hellman key establishment protocols, which can be broken in the quantum computation model. Potential candidates for post quantum key establishment include the ones based on the lattice based Ring Learning With Errors Problem (RING-LWE) introduced in [7,22]. Recently, GOOGLE conducted real life TLS experiments [6] with a RING-LWE based key exchange scheme: the NEWHOPE-USENIX system [6]. While these experiments show the efficiency of NEWHOPE-USENIX, the specification of the reconciliation step of the system is rather complex. The technicality of this step requires a large fraction of the algorithm description in the original paper [1]. This issue together with possible intellectual property right considerations led the designers to introduce a simplified new variant initially named *NewHope-SIMPLE* [25] where the *reconciliation-based* approach of NEWHOPE-USENIX is replaced by an *encryption-based* approach. Thanks to the combined use of *encoding* and *compression* techniques, the performance price to pay for this new version in terms of bandwidth overhead is quite marginal. Now *NewHope-Simple* has been transformed into *NewHope*, a suite of two candidate KEM mechanisms of the NIST call for proposals [23] named *NewHope-CPA-KEM* and *NewHope-CCA-KEM*, in short CPA-KEM and CCA-KEM. Both mechanisms are encryption-based: they rely upon an auxiliary probabilistic public key encryption allowing to encrypt a 256-bit data named CPA-PKE, that is not submitted to the NIST call for proposals as a standalone mechanism.

NewHope-CPA-KEM, that is nearly identical to *NewHope-SIMPLE*, is only claimed to be a passively secure KEM. It can be viewed as the CPA-PKE encryption of a hashed secret random value ν followed by hashing ν on both sides. Unlike CPA-KEM, CCA-KEM is claimed to be secure with respect to adaptively chosen ciphertext attacks. It is derived from CPA-PKE in a less straightforward manner, by applying a variant of the Fujisaki-Okamoto transform [11]. An essential feature of this transform is that the encryption of ν is derandomized: this allows the decrypting party Alice to re-encrypt the decryption result, check that the result matches the received ciphertext, and use this test to prevent information leakages on the private key in active attacks where “dishonestly” derived ciphertext values are sent by an adversary.

While the specification of CPA-KEM and CCA-KEM does not formally prevent re-using the same CPA-PKE (public key, private key) pair in multiple key establishments, the design rationale section of the *NewHope* specification requires

that such key pairs be never cached and that a fresh key pair be generated at each key establishment¹. In the case of CPA-KEM, one of the main reasons for this requirement is that, unlike the classical Diffie-Hellman key establishment, the original RING-LWE based KEM with reconciliation is known to be vulnerable to a practical active attack in a key reuse situation as shown in [10]. Despite not being based on the reconciliation paradigm, CPA-KEM shares sufficiently many features with its predecessor for being conjectured also vulnerable to similar attacks. In the case of CCA-KEM, this requirement to reuse private keys could be justified by the fact that no real perfect forward privacy can be offered if private keys are not ephemeral².

Motivation

With its strong performance and its RING-LWE based security, **NewHope** is a high profile candidate of the NIST competition. There is a good chance for it to be implemented in the future for Internet protocols. So, studying its security under several attacker models is important.

In this paper, we investigate the resilience of the CPA-KEM and CCA-KEM versions of **NewHope** in a misuse situation where the same key pair is reused for multiple key establishment by the private key owner – who will be referred to as Alice in the sequel. Note that Alice is also the party who initiates the two-round key establishment in both schemes. We use the generic name of *key mismatch oracle* to refer to the private key recovery attack models we are considering, that are closely inspired from the adversary model considered in [10]. While slightly less powerful than a CCA attack against an encryption based KEM where a decryption oracle is available, attacks using a *key mismatch oracle* still belong to the active attack category. Their common feature is that the adversary is assumed to be able: (1) to actively interact with Alice by performing multiple KEM establishment where Alice uses the same key pair, (2) to produce each time a guess on the resulting secret key derived by Alice and (3) to access a binary oracle that indicates whether this guess is valid or not.

Our study is motivated by the belief that an in-depth understanding of the security offered by candidate KEM mechanisms submitted to the NIST call for

¹ The single potential exception to this requirement is the *publicseed* part of the public key, whose caching “for say a few hours” seems to be considered by the designers as a viable alternative in situations where the preferred solution of a systematic renewal would turn out to be prohibitively expensive.

² On the other hand this requirement is not fully in line with the former observation, in the **NEWHOPE-USENIX** paper, that “One could enable key caching with a transformation from the CPA-secure key establishment to a CCA-secure key establishment [...]”. Given the performance advantage that may be provided by key caching at server side in certain applications, one can wonder whether it will be strictly followed in practice in all deployments of CCA-KEM if strong cryptanalytic arguments in favour of this conservative choice are not developed during the evaluation of the candidates to the NIST call.

proposals in key reuse situations is a useful part of their cryptanalytic evaluation, even for those candidates for which key reuse is considered as a misuse of the mechanism. Having an accurate estimate of the number of queries to the key mismatch oracle and of the complexity of the private key recovery really helps to assess the possible danger³. We focus here on a case study of the NewHope candidate KEMs. An advantage of this choice is that previous work on reconciliation-based RING-LWE schemes such as [10] can be partly leveraged. However, as will be seen in the sequel, the fact that the NewHope suite is encryption-based and is using encoding techniques induces substantial differences and non-trivially complicates the cryptanalysis. To the best of our knowledge, no investigation of attacks against a scheme without reconciliation in a key mismatch oracle model was published so far.

Previous Work

The danger of accessing a key mismatch oracle within some key agreement protocols in a key share reuse context has been already exposed several times. Early examples showing the vulnerability of some standardized Diffie-Hellman key agreement protocols in such a context were introduced in [20]. The potential danger of a somewhat related type of attack, namely so-called reaction attacks against PKE schemes [13], where an adversary can submit a chosen ciphertext to the legitimate private key owner and access a binary information about her reaction (whether the decryption succeeds or fails for instance), is probably even better known. Bleichenbacher's attack against RSA PKCS#1 of Crypto'98 [4] can be viewed as an early reaction attack. In 1999, Hall, Goldberg and Schneier presented reactions attacks against several PKE schemes [13]. In the particular case of lattice based cryptography, several notes on the vulnerability of NTRU to reaction attacks and its protection against such attacks were published [14, 15]. In 2003, Howgrave-Graham *et al.* proposed a reaction attack on NTRUEncrypt that leverages decryption failures [17]. A recent example of reaction attack is Guo *et al.*'s key recovery attack on the code-based PKE QC-MDPC [12]. It is thus natural that NSA, in 2015, warns NIST Post-Quantum candidates against active attacks [19]. Few times later, the first concrete attacks on a RING-LWE based key establishment leveraging a key mismatch oracle was proposed by Fluhrer [10] (see also [8, 9]).

These attacks rely on the fact that the reconciliation step can be exploited by an active adversary to retrieve some information on the secret static key. Despite the warnings issued in [19], certain NIST candidates are vulnerable to active attacks. Indeed, it is shown in [3] that the secret key of the NIST

³ A similar need to investigate the resilience of candidate algorithms in misuse situations was encountered in the framework of the CAESAR competition aimed at selecting authenticated encryption primitives. In that competition, much analysis was conducted on the resistance of candidates to key recovery attacks in misuse cases such as nonce or decryption-misuse and this provided quite useful information for the algorithms selection process.

candidate HILA5 can be recovered in the **key mismatch oracle** setting following Fluhrer’s approach. In summary, despite the raising awareness of the cryptography research community that **key mismatch oracle** attacks threaten many lattice based KEMs in case of key reuse, relatively few examples of such attacks have been published so far.

About the side channel protection of **NewHope**, no dedicated countermeasure has been proposed for **NewHope** so far, but in [21] a side channel protection for a similar scheme has been proposed. This paper describes a provably first-order secure masking scheme and its integration into a CCA conversion.

Our Contribution

In the following, we evaluate the security of **NewHope** when the attacker gets access to a **key mismatch oracle**. We concretely explain how the attacker can have access to such an oracle in different scenarios with the CPA-KEM and the CCA-KEM. We first introduce a straightforward way to recover such an oracle in the CPA-KEM. The adversary enters a key establishment with Alice, derives from her guess on the shared key produced by Alice a guess on the resulting session key she produces, and attempts to initiate a session with Alice under this guessed session key. The success or failure of this protected communication attempt provides the desired **key mismatch oracle**⁴.

Then, at the end of the paper, we elaborate other scenarios on the CCA version which require side channels. Indeed, the CCA-KEM version induces major extra differences with formerly analyzed reconciliation-based schemes, that also deserve being analyzed. Because of the CCA transform, a **key mismatch oracle** cannot be accessed directly. But we show that for insufficiently protected implementations, simple faults or side channels could bypass this transform and provide the desired **key mismatch oracle**. While unprotected versions of CCA-KEM are extremely efficient, its implementations must be very carefully protected against any **key mismatch oracle** leakage if key pairs are potentially reused. This might eventually come with a cost in terms of performance. This study may help the developers to protect the algorithms against a possible **key mismatch oracle** leakage.

The core of this work is the description of a new attack on **NewHope** using the **key mismatch oracle**. Even if the existence of previous work attacks (see [8, 10]) casts suspicion on the resistance of **NewHope** CPA-KEM against active attacks in the same key-reuse setting, one has to take into account substantial differences between the reconciliation-based paradigm of the original **NewHope** and the encryption-based paradigm of CPA-KEM. Because of these differences, the detail of Fluhrer’s attack [10] is not really inspiring for mounting an attack and any direct transposition attempt would be hopeless. Finding an efficient way

⁴ It is worth noticing that the same direct access to a **key mismatch oracle** remains feasible if the KEM exchange is embedded in an authenticated key establishment protocol, under the sole condition that the adversary is the owner of a valid authentication (or signature) key.

of deriving information on the secret key from the key mismatch oracle with a low number of queries induces several issues. The main difficulty is to retrieve enough leakages after the application of the encoding and compression functions. We investigated how to leverage these functions in order to find a simple way to instantiate the oracle and identified precise elements in the polynomial ring that can be used by the adversary to recover the secret. Finally, we had to take into account the fact that **NewHope** coefficients are in $[-8, 8]$.

We experimented our attack with a **Magma** CAS proof of concept. Under **NewHope** parameters, we were able to recover exactly the secret **S** with on average 16,700 queries for **NewHope1024** which corroborates the expected performance of the model.

Paper Outline

In Sect. 2 we introduce some notation and describe the **NewHope** CPA scheme. In Sect. 3, we describe the notion of key mismatch oracle and how practical it can be for the CPA-KEM. In Sect. 4, we detail our attack using the key mismatch oracle. In Sect. 4.4, we present our experiments. In Sect. 5, we show how the key mismatch oracle can be retrieved with side channels with the CCA-KEM. Finally, in Sect. 6, we summarize our results and discuss future research.

2 Preliminaries

2.1 Notations

Let q be a prime number in \mathbb{N} and let \mathbb{Z}_q denote the ring elements $\mathbb{Z}/q\mathbb{Z}$. Depending on the context, the elements in \mathbb{Z}_q can be equivalently represented as integers in $\{0, \dots, q-1\}$ or in $\{-\frac{q-1}{2}, \dots, \frac{q-1}{2}\}$. In the following, the notation \mathcal{R} refers to the polynomial ring $\mathbb{Z}_q[x]/(x^N + 1)$ with N a power of 2. If \mathbf{P} belongs to \mathcal{R} , it is a polynomial of degree $(N - 1)$ with coefficients $\mathbf{P}[i]$ belonging to the set \mathbb{Z}_q . Such elements can also be represented as vectors whose i -th coordinate is the coefficient related to x^i . In the sequel we use either the polynomial notation or the vectorial one. For readability, bold capital letters are used to refer to elements in \mathcal{R} and bold lowercase letters will refer to compressed elements, i.e. elements in \mathcal{R} with small coefficients.

Let us define \mathcal{G}_a as the centered Gaussian distribution of standard deviation a and ψ_k the centered binomial distribution of parameter k . Its standard deviation is $\sqrt{k}/2$. One may sample from ψ_k for integer $k > 0$ by computing $\sum_{i=1}^k b_i - b'_i$, where the $b_i, b'_i \in \{0, 1\}$ are uniform independent bits.

Property 1. The elements generated according to a centered binomial distribution ψ_k of parameter k are in the interval $[-k, k]$. Thus, the coefficients of the small elements drawn from \mathcal{R} in **NewHope** are in $[-8, 8]$.

In the figures and algorithms, the notation $\stackrel{\$}{\leftarrow} \mathcal{D}$ means picking an element in \mathcal{R} having all its coefficients generated at random according to distribution \mathcal{D} . The notation $\stackrel{coin}{\leftarrow} \mathcal{D}$ means using a *coin* $\in \{0, \dots, 255\}$ ³² as a seed to pick a pseudorandom element in \mathcal{R} having all its coefficients according to distribution \mathcal{D} . This is generally done using a hash function like SHAKE-128. In the paper we refer several times to $Sign(a)$ with $a \in \mathbb{Z}$ by using the convention that it is defined as *positive* when $a \geq 0$ and as *negative* when $a < 0$. If $x \in \mathbb{R}$, the integer $\lfloor x \rfloor$ is defined as $\lfloor x + \frac{1}{2} \rfloor \in \mathbb{Z}$.

2.2 NewHope

NewHope [23, 25] is a RING-LWE based key establishment scheme derived from NEWHOPE-USENIX [1], that is simplified because it does not use the reconciliation anymore. In this section, we describe NewHope, where we omit some details (e.g. the so-called NTT transform or the encoding of the messages) to simplify the presentation. This does not imply any loss of generality for our attack. To ease the understanding, we will describe the CPA-KEM version of NewHope in this section as the key mismatch oracle can be easily derived. We will present the CCA-KEM later in Sect. 5 when we present some ways to access a key mismatch oracle.

The polynomial ring \mathcal{R} used in NewHope has the following parameters: $(N, q) = (1024, 12289)$ or $(N, q) = (512, 12289)$. The coefficients of the small elements drawn from \mathcal{R} follow a centered binomial distribution ψ_k^N with $k = 8$. The standard deviation is $a = \sqrt{\frac{8}{2}} = 2$. We decided to focus on explaining the attack for $N = 1024$. Indeed, for $N = 512$ there is twice less redundancy and the attack is easier. Thus, we fix $N = 1024$. These elements will be seen as vectors of size N with integer components. We denote $s = 1536$ which is such that $q = 8s + 1$. The aim of the system is to share a key of size 256 bits following the exchange mechanism outlined below and represented in Fig. 1.

A public value $\mathbf{A} \in \mathcal{R}$ is derived from a published *seed*. Four specific functions are introduced: Encode, Decode, Compress and Decompress. They are described in Algorithms 1, 2, 3 and 4. Note that we partly deviate from the notation of the original specification of these algorithms, since we use the parameter s (the original description is in [25]). The following paragraphs describe these functions.

Compress and Decompress. The function Compress (Algorithm 3) takes as input a vector \mathbf{C} in \mathcal{R} and applies on each of its component a *modulus switching* to obtain an element \mathbf{c} in $\mathbb{Z}_8[x]/(x^N + 1)$. Compressing a vector \mathbf{C} essentially means keeping the 3 most significant bits of each coefficient. The function Decompress (Algorithm 4) shifts the bits of the input $\mathbf{c} \in [0, 8]^N$ to place them among the most significant bits. These functions are not the inverse of each other.

Encode and Decode. The Encode function takes a n -bit input ν where $n = N/4$ and creates an element $\mathbf{K} \in \mathcal{R}$ which stores 4 times the element ν . The redundancy is used by the function Decode to recover ν with a noisy \mathbf{K} .

NewHope Key Encapsulation Mechanism. Let us now describe this scheme.




Fig. 1. Simplified NewHope

1. *Setup*: Alice generates 2 small secrets \mathbf{S} and \mathbf{E} in \mathcal{R} . She sends $\mathbf{B} = \mathbf{AS} + \mathbf{E}$ to Bob.
2. *Key Encapsulation*: From a random *coin* acting as a seed, Bob derives 3 small secrets \mathbf{S}' , \mathbf{E}' and \mathbf{E}'' in \mathcal{R} and a random element ν_B of size n which will be the encapsulated key. He computes $\mathbf{U} = \mathbf{AS}' + \mathbf{E}'$. He encodes ν_B into a redundant element \mathbf{K} of \mathcal{R} using the algorithm `Encode` (Algorithm 1). Bob uses `Compress` (Algorithm 3) to compress $\mathbf{C} = \mathbf{BS}' + \mathbf{E}'' + \mathbf{K}$ into an element with very small coefficients as described above. He sends $(\mathbf{c} = \text{Compress}(\mathbf{C}), \mathbf{U})$ to Alice. He deduces the shared secret as $\mu_B = \text{SHAKE-256}(32, \nu_B)$.
3. *Key Decapsulation*: Alice decompresses \mathbf{c} with `Decompress` into \mathbf{C}' (Algorithm 4). She computes $\mathbf{C}' - \mathbf{US}$ which is close to

$$\mathbf{C} - \mathbf{US} = \mathbf{ES}' + \mathbf{E}'' + \mathbf{K} - \mathbf{E}'\mathbf{S}. \quad (1)$$

Algorithm 1. Key Encoding

```

1 function Encode( $\nu \in \{0, 1\}^n$ )
2  $\mathbf{k} \leftarrow 0$ 
3 for  $i := 0$  to  $n - 1$  do
4      $\mathbf{K}_i \leftarrow \nu_i \cdot 4s$ 
5      $\mathbf{K}_{i+n} \leftarrow \nu_i \cdot 4s$ 
6      $\mathbf{K}_{i+2n} \leftarrow \nu_i \cdot 4s$ 
7      $\mathbf{K}_{i+3n} \leftarrow \nu_i \cdot 4s$ 
8 end
9 Return  $\mathbf{k}$ 
    
```

Algorithm 2. Key Decoding

```

1 function Decode( $\mathbf{K} \in \mathcal{R}$ )
2  $\nu \leftarrow 0$ 
3 for  $i := 0$  to  $n - 1$  do
4      $t \leftarrow \sum_{j=0}^3 \mathbf{K}_{i+jn} - 4s$ 
5     if  $t < q$  then  $\nu_i \leftarrow 1$  else  $\nu_i \leftarrow 0$ 
6 end
7 Return  $\nu$ 
    
```

Algorithm 3. Compression

```

1 function Compress( $C \in \mathcal{R}$ )
2 for  $i := 0$  to  $N - 1$  do
3    $c[i] \leftarrow \left\lceil \frac{8 \cdot C[i]}{q} \right\rceil \bmod 8$ 
4 end
5 Return  $c$ 

```

Algorithm 4. Decompression

```

1 function Decompress( $c \in [0, 8]^N$ )
2 for  $i := 0$  to  $N - 1$  do
3    $C'[i] \leftarrow \left\lceil \frac{q \cdot c[i]}{8} \right\rceil$ 
4 end
5 Return  $C'$ 

```

Since $\mathbf{ES}' + \mathbf{E}'' - \mathbf{E}'\mathbf{S}$ is small, she recovers an estimated value ν_A of ν_B with a decoding algorithm called Decode(Algorithm 2). From ν_A , she can deduce $\mu_A = \text{SHAKE-256}(32, \nu_A)$.

Since $\mathbf{S}, \mathbf{S}', \mathbf{E}, \mathbf{E}', \mathbf{E}''$ are small, Alice and Bob get the same key $\mu = \mu_B = \mu_A$ with high probability.

Remark 1. This Section presented NewHope-CPA-KEM which is the target Sect. 4’s analysis. However a PKE called NewHope-CPA-PKE has been introduced in [23]. The slim difference lies on the fact that ν_B becomes the encrypted message. The CCA security of the CCA version, called NewHope-CCA-KEM relies on the CPA security of NewHope-CPA-PKE (see Sect. 5).

3 The Key Mismatch Oracle

This section introduces the notion of key mismatch oracle and a way to access it in the CPA version. We will always consider a malicious active adversary, Eve, who acts as Bob. Her messages, key and intermediate values will be denoted as m_E, μ_E and ν_E instead of m_B, μ_B and ν_B .

Remark 2. One might wonder how a malicious Alice can recover Bob’s secret in a case of key reuse by Bob. In NewHope, this can be done with 2 queries, see the full version of our paper [2].

The goal of the adversary is to recover Alice’s static private keys \mathbf{S} and \mathbf{E} by using the following oracle several times. We will focus on recovering the secret \mathbf{S} . \mathbf{E} can be derived from \mathbf{S} with $\mathbf{E} = \mathbf{B} - \mathbf{AS}$.

Definition 1 (*key mismatch oracle*). A key mismatch oracle is an oracle that outputs a bit of information on the possible mismatch at the end of the key encapsulation mechanism.

In the NewHope context, the key mismatch oracle is the oracle that takes any message m_E and any key hypothesis μ_E as input and outputs the following

$$\mathcal{O}_1(m_E, \mu_E) = \begin{cases} 1 & \text{if } (\mu_A =) \text{Decapsulation}(m_E, \mathbf{S}) = \mu_E \\ -1 & \text{otherwise} \end{cases} \quad (2)$$

Such an oracle should leak information on secret \mathbf{S} because its output is clearly correlated to the value of \mathbf{S} . However, this oracle is less powerful than a CCA decryption oracle against CPA-PKE. Indeed, the only information given is a bit representing the possible key mismatch. The difficulty is to choose appropriate (m_E, μ_E) to retrieve information of a small part of \mathbf{S} . In Sect. 4, we present how to recover the secret \mathbf{S} from such an oracle.

The simplest way to access such an oracle is when the CPA-KEM is implemented with static secrets. In other words, Alice will keep her secrets \mathbf{S} and \mathbf{E} for several key establishment requests. We consider that Eve does not necessarily follow the scheme specification. She can “cheat” and generate a message m_E that is not derived from a coin or from random small secrets \mathbf{S}' , \mathbf{E}' and \mathbf{E}'' . By definition, the CPA version of NewHope is passively secure, an attacker using a key mismatch oracle is outside of the security assumptions⁵. This has been well highlighted in paragraph 2.3, Section **No key caching** of the original paper of NEWHOPE-USENIX. However an implementation of NewHope which allows misuse cases (see [24]) cannot be completely excluded. Thus it is important to precisely evaluate such a threat and consider the following attack model.

Attack Model 1. *Alice will accept any syntactically correct message m_E and always try to use the corresponding shared key for communicating. When she derives the shared key, either she is able to decrypt messages exchanged after that with Eve (and thus Eve deduces that the shared key is the same) or she will notify Eve that something went wrong with the key agreement. Eve will then deduce that the key is different. In both cases, Eve gets the desired key mismatch oracle.*

In Sect. 5, we show how to get access to such an oracle with side channels in the CCA framework.

4 Attack on NewHope with Key Mismatch Oracle

We assume here that Eve, the attacker, has access to \mathcal{O}_1 , a key mismatch oracle as defined in Sect. 3. Let us now explain how she proceeds to recover Alice’s static secret key \mathbf{S} following Attack Model 1.

⁵ While key reuse is against the designers’ requirements of the NIST submission NewHope, as expressed in the footnote in the design rationale on p. 16, this requirement does not seem to be formally reflected in the algorithm description of Sect. 1.2. This section indeed defines separate algorithms for key pairs generation, (en/de)apsulation, but does not state that a pair shall be used only once. Thus, though running NewHope with key reuse represents a misuse situation, analyzing the security of this scheme in this situation is definitely much more relevant question than considering variations in the formal specification of NewHope and investigating resulting weaknesses.

4.1 Rewriting the Key Mismatch Oracle

The use of the key mismatch oracle obviously leaks information on Alice’s secret key \mathbf{S} . But the task of recovering \mathbf{S} entirely seems much more complicated. Indeed as defined in Sect. 3, the only information provided by the key mismatch oracle is a bit representing the success or mismatch of the key agreement. The difficulty for Eve is to choose appropriate (m_E, μ_E) pairs to get useful information on small parts of \mathbf{S} .

In a first step, Eve simplifies her part of the protocol in such a way that the knowledge of the key mismatch oracle output bit $\mathcal{O}_1(m_E, \mu_E)$ can be easily exploited. To do so, she can fix for instance μ_E such that:

$$\nu_E = (1, 0, \dots, 0) \text{ and thus } \mu_E = \text{SHAKE-256}(32, \nu_E). \tag{3}$$

The value of $\nu_E = (1, 0, \dots, 0)$ has not been arbitrarily chosen; as we will see later, the 0 in positions 1 to $n - 1$ will help the success rate of the attack (see Proposition 3). From now on, the value of μ_E is *fixed* according to Eq. (3). Moreover, when replacing m_E by its definition: $m_E = (\mathbf{c}, \mathbf{U})$, the oracle \mathcal{O}_1 can be reformulated using the oracle \mathcal{O}_2 defined below.

Definition 2. *Let us introduce oracle \mathcal{O}_2 such that $\mathcal{O}_2(\mathbf{c}, \mathbf{U}) = \mathcal{O}_1((\mathbf{c}, \mathbf{U}), \mu_E)$.*

With this new definition, Eve can adapt the values of \mathbf{c} and \mathbf{U} to leverage Oracle \mathcal{O}_2 and retrieve information on \mathbf{S} . In other words, since μ_E is fixed, the inputs (\mathbf{c}, \mathbf{U}) are the degrees of liberty for finding \mathbf{S} .

From Alice’s side, the link between ν_A and \mathbf{S} passes through the functions Decode, Decompress (see the figures in full version of our paper [2].) and the element \mathbf{K}' : $\nu_A = \text{Decode}(\mathbf{k}') = \text{Decode}(\mathbf{C} - \mathbf{U}\mathbf{S}) = \text{Decode}(\text{Decompress}(\mathbf{c}) - \mathbf{U}\mathbf{S})$. Thus, from the definition of the Decode algorithm, the value of $\nu_A[i]$, the i -th component of ν_A , is deduced from the following sign computation:

$$\text{Sign} \left(\sum_{j=0}^3 \left| (\text{Decompress}(\mathbf{c}) - \mathbf{U} \cdot \mathbf{S})[i + nj] - 4s \right| - q \right) \tag{4}$$

We recall here that 0 is positive by convention.

The problem for Eve is that she is unable to know the number of errors that will occur at the end of the decryption computations and the positions in which they appear. Indeed, the key mismatch oracle only gives one bit of information corresponding either to *mismatch* or *success*. If there is a mismatch, Eve knows that *at least* one bit of ν_A is different from ν_E but she can not determine which one (or which ones). Therefore, in order to mount an effective attack, Eve needs to restrict all these different possibilities by making the following hypothesis:

Hypothesis 1. *For i from 1 to $n - 1$, the component $\nu_A[i]$ is equal to 0.*

If Hypothesis 1 is verified, any failure in the communication comes from a single error in ν_A located in the very first component $\nu_A[0]$. Indeed, in that case, the success of the exchange only depends on the first computed value $\nu_A[0]$. In particular, if we assume this hypothesis, the oracle \mathcal{O}_2 depends only on the $\nu_A[0]$ and we obtain the following result.

Lemma 1. *Under Hypothesis 1, the initial oracle \mathcal{O}_2 can be rewritten as*

$$\mathcal{O}_2(\mathbf{c}, \mathbf{U}) = \text{Sign} \left(\sum_{j=0}^{j=3} \left| (\text{Decompress}(\mathbf{c}) - \mathbf{U}\mathbf{S})[0 + nj] - 4s \right| -q \right)$$

For mounting her attack, Eve has to find pairs (\mathbf{c}, \mathbf{U}) that

1. target the smallest number of bits of \mathbf{S}
2. verify Hypothesis 1

For item 1, since the Decode algorithm takes coefficients of \mathbf{S} four by four, the size of the smallest target is a quadruplet of coefficients of \mathbf{S} . Actually, for a given quadruplet of integers $\underline{\ell} = (\ell_0, \ell_1, \ell_2, \ell_3)$ and a target index k (i.e. an index corresponding to the components of \mathbf{S} that Eve wants to retrieve), by taking

$$\mathbf{U} = sx^{-k} \quad \text{and} \quad \mathbf{c} = \sum_{j=0}^3 \left((\ell_j + 4) \bmod 8 \right) \cdot x^{nj} \quad (5)$$

one can prove (see in Proposition 1) that Eve targets the quadruplet $(\mathbf{S}[k + nj])_{j=0\dots 3}$. Indeed, the element x^{-k} will “rotate” \mathbf{S} in order to target $(\mathbf{S}[k + nj])_{j=0\dots 3}$ and \mathbf{c} is induced by the quadruplet $\underline{\ell} = (\ell_0, \ell_1, \ell_2, \ell_3)$ that can vary.

About item 2, with this choice of (\mathbf{c}, \mathbf{U}) , the Hypothesis 1 has good chances to be verified because the coefficients of \mathbf{c} outside from the set $\{k + nj \mid j = 0 \dots 3\}$ are 0. So, the same coefficients of $\mathbf{C} - \mathbf{U}\mathbf{S}$ have good chances to be small. Then, Alice is likely to derive 0 for these coefficients of ν_A . However, it is not always verified and this will impact the attack’s success rate. We will discuss and compute this probability later in Proposition 3.

We can now introduce \mathcal{O}_3 , a reformulation of \mathcal{O}_2 depending on target index k and the quadruplet $\underline{\ell}$ (see Eq. 5):

$$\mathcal{O}_3(k, \underline{\ell}) = \mathcal{O}_2 \left(sx^{-k}, \sum_{j=0}^3 \left((\ell_j + 4) \bmod 8 \right) \cdot x^{nj} \right)$$

This formulation of the key mismatch oracle is more convenient in order to explain how Eve will gather information on \mathbf{S} from instantiations of $\underline{\ell}$. The following proposition shows a first result in this direction.

Proposition 1. Final oracle. *Let us assume that Hypothesis 1 is verified. Let k be a target index ($k \in [0, n - 1]$). For a given integer quadruplet $\underline{\ell}$ in $[-4, 3]^4$, the (\mathbf{c}, \mathbf{U}) explicited in Eq. 5 is such that*

$$\mathcal{O}_3(k, \underline{\ell}) = \text{Sign} \left(\sum_{j=0}^{j=3} \left| \ell_j - \mathbf{S}[k + nj] \right| -8 \right)$$

Proof. The proof is given in the full version of our paper [2].

In the next section, we explain how to effectively use the form \mathcal{O}_3 of the key mismatch oracle to extract the secret \mathbf{S} .

4.2 Recovering Very Small Coefficients of \mathbf{S}

Let us recall that the secret \mathbf{S} is a polynomial in $\mathbb{Z}_q[x]/(x^N + 1)$ with coefficients in $[-8, 8]$, it can be seen as a vector of N components $\mathbf{S}[i]$. Eve will recover the coefficients of the secret \mathbf{S} four by four. Let k be the index of the targeted quadruplet $[\mathbf{S}[k], \mathbf{S}[k + n], \mathbf{S}[k + 2n], \mathbf{S}[k + 3n]]$. The index k goes from 0 to $n - 1$ and for each fixed k , Eve will call the oracle $\mathcal{O}_3(k, \ell)$ with several appropriate value of ℓ until she gets the secret values.

For simplicity, let us now **fix the index** k and denote $S_j = \mathbf{S}[k + nj]$.

The following proposition and corollary describe an algorithm that, when iterated (see Corollary 1), allows to recover S_j for j from 0 to 3.

Proposition 2. *Let us fix j in $[0, 3]$. Under Hypothesis 1, if S_j is in $[-3, 2]$ and $(S_i)_{i \neq j} \in [-4, 4]$, there exists a probabilistic algorithm \mathcal{A} which recovers the value S_j in 8 queries to oracle \mathcal{O}_3 with a success probability depending on the distribution of $(S_i)_{0 \leq i \leq 3}$.*

Corollary 1. *Under Hypothesis 1, if S_j is in $[-3, 2]$ and $(S_i)_{i \neq j} \in [-4, 4]$, there exists a probabilistic algorithm \mathcal{A}' which recovers the value S_j with an average number of queries to oracle \mathcal{O}_3 depending on the distribution of $(S_i)_{0 \leq i \leq 3}$.*

In the sequel of this section, we give the proof of Proposition 2 by first presenting the construction of the algorithm and then by introducing a method to assess the success rate. We refer the reader to the full version of our paper [2] for the proof of Corollary 1.

Proof of Proposition 2

Description of \mathcal{A} . Let us prove the proposition by focusing on the secret S_0 and by explaining how it can be recovered in 8 queries to oracle \mathcal{O}_3 . The process will then be exactly the same for the three other values S_1, S_2 and S_3 .

The first step consists in taking the 3 values ℓ_1, ℓ_2, ℓ_3 at random inside the interval $[-4, 3]$. Knowing that all S_j are fixed, the quantity $\sum_{j=0}^3 |\ell_j - S_j| - 8$ can thus be expressed by $f_v(\ell_0) = |\ell_0 - S_0| + v - 8$ with $v = \sum_{i=1}^3 |\ell_i - S_i|$ a fixed unknown constant (since all S_j are unknown). Let us now see how $f_v(\ell_0)$ behaves when ℓ_0 varies, see Fig. 2 for an illustration.

We now assume that one makes 8 queries to the oracle \mathcal{O}_3 : one for each value of ℓ_0 varying inside $[-4, 3]$. Such queries imply having access to $Sign(f_v(\ell_0)) \forall \ell_0 \in [-4, 3]$. The analysis can thus be split in 2 cases:

1. If $(v - 8) \geq 0$ then all queries to oracle 1 obviously lead to “positive signs”. It is quite clear when one looks at Fig. 2.
2. If $(v - 8) < 0$, two subcases occur
 - In some cases, there exists two possible values $\tau_1 < \tau_2$ such that the function $|\ell_0 - S_0| + (v - 8)$ goes from a positive value to a negative one at point τ_1 and then from a negative value to a positive one at point τ_2 . We call this case the *favorable case*. Figure 3 provides a good illustration.




Fig. 2. If $v - 8 \geq 0$




Fig. 3. If $v - 8 > 0$




Fig. 4. If $v - 8 \gg 0$

$$\frac{\ell_0 - 4 \cdots \tau_1 - 1 \quad \tau_1 \tau_1 + 1 \cdots \tau_2 - 1 \quad \tau_2 \tau_2 + 1 \cdots 3}{0 + \cdots + \quad + \quad - \quad \cdots - \quad + + \quad \cdots +}$$

- If $(v - 8) < 0$ and $v \ll 8$, only one change of sign will occur in the interval $[-4, 3]$. Figure 4 provides a good illustration.

Figure 3 illustrated what happens in the *favorable case*. Around S_0 , the trace has a slope equal to $+$ or -1 . Because of the symmetry, the value S_0 can simply be recovered by:

$$S_0 = \frac{\tau_2 + \tau_1}{2}. \tag{6}$$

If we are not in the *favorable case*, two such values τ_1 and τ_2 do not exist. This means that the constant v is not appropriate.

Termination of \mathcal{A} . For any $S_0 \in [-3, 2]$, \mathcal{A} has a non zero success probability. Indeed, no matter the values of (S_1, S_2, S_3) in $[-4, 4]^3$, the 3-uple $(\ell_1, \ell_2, \ell_3) \in [-4, 3]^3$ defined by

$$\left(\ell_1 = S_1 - 2 \cdot \text{Sign}(S_1), \ell_2 = S_2 - 2 \cdot \text{Sign}(S_2), \ell_3 = S_3 - 3 \cdot \text{Sign}(S_3) \right)$$

is at least one of the choices inducing a favorable case. Actually, one can check that this choice implies that $v = 7$. Thus $v - 8 = -1$ which always gives a favorable case for finding $S_0 \in [-3, 2]$.

Table 1. Success probability of \mathcal{A} for $(S_j)_{1 \leq j \leq 3}$ following ψ_4 distribution

| | | | | | | |
|--|-----|-----|-----|-----|-----|-----|
| S_0 | -3 | -2 | -1 | 0 | 1 | 2 |
| Probability (%) | 14 | 27 | 39 | 39 | 27 | 14 |
| Expected number of iterations ($1/\text{probability}$) | 7.1 | 3.7 | 2.6 | 2.6 | 3.7 | 7.1 |

Success Probability. A precise probabilistic study on the $(S_j)_{1 \leq j \leq 3}$ to assess the success rate of algorithm \mathcal{A} is detailed in the full version of our paper [2]. In Table 1, one can find the probability of success assuming that S_1, S_2, S_3 follow a binomial distribution ψ_4 . The expected number of iterations is the average amount of tries before recovering the secret. \square

Example 1. Let us suppose that $S_i = [0, -2, 1, -1]$. For $(\ell_1, \ell_2, \ell_3) = (2, -2, -1)$, $\sum_{j=0}^3 |\ell_j - S_j| - 8 = |\ell_0 - S_0| - 2$. If we query the sign of the latter for $\ell_0 = -4, -3, -2, -1, 0, 1, 2, 3$, we get: +, +, +, -, -, -, +, +. We can conclude that $S_0 = \frac{1-1}{2} = 0$. Whereas, for $(\ell_1, \ell_2, \ell_3) = (-2, 0, 1)$, $\sum_{j=0}^3 |\ell_j - S_j| - 8 = |\ell_0 - S_0| - 5$. The sign for $\ell_0 = -4, -3, -2, -1, 0, 1, 2, 3$ becomes: -, -, -, -, -, -, -. We cannot conclude anything on S_0 .

At the end of this section, with Corollary 1, we know that if \mathbf{S} is generated with coefficients following the ψ_4 distribution and if Hypothesis 1 is verified, there exist an algorithm that recovers each coefficient of \mathbf{S} that is in $[-3, 2]$ (i.e. almost 96% of the coefficients). If a coefficient of \mathbf{S} is not in $[-3, 2]$, no favorable case will appear and the coefficient will not be found. In the next section, we adapt this method for NewHope.

4.3 Recovering \mathbf{S} for NewHope Parameters

In this section, we describe a way to recover \mathbf{S} for NewHope parameters, i.e. when the binomial parameter is 8. According to Property 1, the coefficients of $\mathbf{S}[k]$ are in $[-8, 8]$. This is outside from the hypothesis made in Proposition 2. Indeed, the coefficients $\mathbf{S}[k]$ should lie in $[-3, 2]$. One can make the following change in order to fit with Proposition 2 hypothesis: $\mathbf{S}^1 = \frac{\mathbf{S}}{2}$. In order to target \mathbf{S}^1 instead of \mathbf{S} , one can change \mathbf{U} from Eq. 5 to be the following $\mathbf{U} = \frac{s}{2}x^{-k}$.

Let us wrap up the attack into the following Proposition.

Proposition 3. *There exists a probabilistic algorithm \mathcal{B} which recovers NewHope secret \mathbf{S} with high probability using an average of 18,500 queries for $N = 1024$.*

Proof. Let $k \in [0, n - 1]$. The distribution of probabilities for $\mathbf{S}[k]$ is in Table 2.

Case 1: $\mathbf{S}[k]$ belongs to $\{-8, -7, 5, 6, 7, 8\}$. The probability of this case is around 1%. In that case, at most one change of sign will always happen. Then, only the sign of $\mathbf{S}[k]$ can be recovered and a brute force should be done at the end of the attack to distinguish among the possible values. For $N = 1024$, on average 10 coefficients out of 1024 will not be found. When a positive value is not found, it has 8/10 chances to be a 5. At the end of the attack, a brute-force step evaluating $\mathbf{B} - \mathbf{AS}$ and taking account of the probabilities can be done.

Table 2. Distribution ψ_8 (note that the probability is the same for negative values)

| $\mathbf{S}[k]$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|---------------------------------|-------|-------|------|------|------|-----|-----|----|---|
| Probability ($\times 2^{16}$) | 12870 | 11440 | 8008 | 4368 | 1820 | 560 | 120 | 16 | 1 |




Fig. 5. When $S[k] \bmod 2 = 0$




Fig. 6. When $S[k] \bmod 3 = 1$

Case 2: $S[k]$ belongs to $\{-6, \dots, 4\}$. In that case, $S^1[k]$ belongs in the interval $[-3, 2]$. The attack is the one from Proposition 2 with a different secret $S^1[k] = \frac{S[k]}{2}$. However, the results will not be as accurate as before. We will show that there is a subtlety that allows Eve to recover the exact value of $S[k]$. There are 2 possible results depending on $S[k] \bmod 2$:

- If $S[k] \bmod 2 = 0$, then $S^1[k] \in \{-3, -2, -1, 0, 1, 2\}$. Proposition 2 allows Eve to recover S^1 . In other words, Eve will recover a succession of signs where an odd number of $(-)$ occurs (see Fig. 5). She will then be able to recover $\frac{S[k]}{2} = \frac{\tau_1 + \tau_2}{2}$ and then $S[k]$.
- If $S[k] \bmod 2 = 1$ then $S^1[k] \in \{-2.5, -1.5, -0.5, 0.5, 1.5\}$. In a favorable case of Proposition 2, the situation will be different. As in Fig. 6, the number of $(-)$ is then even.

Wrap Up. Here is a procedure to recover $S[k]$.

- Case 1.* If the number of $(-)$ is odd, then $S[k]$ is even and $S[k] = 2 \frac{\tau_1 + \tau_2}{2} = \tau_1 + \tau_2$
- Case 2.* If the number of $(-)$ is even, then $S[k] = 2 \left\lfloor \frac{\tau_1 + \tau_2}{2} \right\rfloor + 1$
- Case 3.* If at most one change of sign occur, the procedure is restarted.

If the number of restarts is too large (say, $\geq M$), the procedure is stopped and the coefficient, placed in a bruteforce set, is found at the end of the attack.

Table 3. Average number of queries

| Value | -6 | -5 | -4 | -3 | -2 | -1 | 0 | 1 | 2 | 3 | 4 | -8, -7, 5, 6, 7 or 8 |
|-----------------|----|----|----|----|----|----|----|----|----|----|----|----------------------|
| Average queries | 33 | 33 | 19 | 20 | 16 | 17 | 17 | 15 | 22 | 20 | 38 | M |

Number of Queries. The amount of queries is derived with the same technique as in the full version of our paper [2]. See Table 3 for the average number of queries. Let us set the threshold M to 50, to get the total average number of queries, we compute the expected number of queries for $S[k]$ (≈ 18) and multiply it by $N = 1024$.

Success Probability. The success probability depends only on Hypothesis 1 with \mathbf{S}^1 , which becomes the following for \mathbf{S} .

Hypothesis 2. $\forall i, k \in [1, n - 1] \sum_{j=0}^{j=3} \left| \frac{\mathbf{S}[k+i+nj \bmod N]}{2} + 4 \right| \geq 8$

Hypothesis 2 is true with a probability 94.6% for $N = 1024$. Indeed, to compute this probability, one can check whether each quadruplet verifies it. Only a few unlikely quadruplet (e.g. [8, 8, 8, 8]) do not verify the hypothesis. \square

4.4 Experimental Results

We implemented a proof of concept with Magma CAS [5]⁶. We coded NewHope according to its parameters and used the key mismatch oracle for the attack. We worked on a basic optimization of the number of queries. We ran 1000 experiments and recovered more than 95% of the secret keys in an average time of 30 minutes per key and 16,700 queries. We still think that the number of queries and the time can be better optimized.

5 Accessing the Key Mismatch Oracle with the CCA Version of NewHope

In order to be protected against active attacks, the CPA-KEM of NewHope has been transformed according to the Hofheinz, Hövelmanns and Kiltz CCA transformation [16] which is a variant of the Fujisaki-Okamoto transformation [11]. The CCA security is then based on the CPA security of the PKE. The CCA transformation of the algorithms defining this version of NewHope is detailed in Algorithms 5, 6 and 7. These algorithms use the underlying CPA-PKE of NewHope as defined in Sect. 1.2.1 of [23].

Algorithm 5. NewHope CCA-KEM Key Generation

```

1 function NewHope-CCA-KEM.Gen()
2 (pk, sk) ← NewHope-CPA-PKE.Gen()
3 s ← {0, ..., 255}32
4 return (pk, sk = (sk||pk||SHAKE-256(32, pk)||s))

```

Algorithm 6. NewHope-CCA-KEM Encapsulation

```

1 function NewHope-CCA-KEM-Encaps(pk)
2 coin ← {0, ..., 255}32
3 μ ← SHAKE-256(32, coin) ∈ {0, ..., 255}32
4 K||coin' ||d ←
   SHAKE-256(96, μ||SHAKE-256(32, pk)) ∈
   {0, ..., 255}32+32+32
5 c ← NewHope-CPA-PKE.Encrypt(pk, μ; coin')
6 ss ← SHAKE-256(32, K||SHAKE-256(32, c||d))
7 return (c̄ = c||d, ss)

```

One can note the main security measure in Algorithm 7 where the instruction in red corresponds to a double encryption to check if the message m_B has been honestly generated.

⁶ The Magma code can be found at <https://www.di.ens.fr/~mrossi/>.

More precisely, in the key mismatch oracle, the message m_B can be adjusted by the attacker but with this CCA version of **NewHope**, Eve must follow the protocol and generate m_E according to a seed called $coin'$ that is derived from μ_E and another seed called $coin$ (a 32-byte random integer). Then, Alice will derive $coin'$ to check if μ_E was computed following the protocol. Then, a key mismatch will come from the following oracle

$$\mathcal{O}_4(coin, \mu_E) = \begin{cases} 1 & \text{if } \text{Dec}(\text{Enc}(m_A, coin), \mathbf{S}) = \mu_E \\ -1 & \text{otherwise} \end{cases} \quad (7)$$

This oracle is less convenient than \mathcal{O}_1 because with an honest behaviour, the error probability is claimed to be lower than 2^{-213} in the NIST specification (paragraph 4.2.7 of [23]). In the sequel, we point at critical steps inside the CCA transform that let Eve access oracle \mathcal{O}_1 using side channel or fault attacks.

On Using Side Channel or Fault Attack

When the attacker has access to a device implementing Alice's side of the exchange, the attack model should take into account situations where some algorithmic security measures may be bypassed by using hardware attacks.

Power Analysis. Here we consider that Eve is able to make a power analysis during the verification step of Alice decapsulation algorithm.

Algorithm 7. NewHope-CCA-KEM Decapsulation

```

1 function NewHope-CCA-KEM.Decaps( $\bar{c}, \bar{s}k$ )
2  $c||d \leftarrow \bar{c} \in \{0, \dots, 255\}^{\text{NEWHOPE\_CPAPKE\_CIPHERTEXTBYTES}+32}$ 
3  $sk||pk||h||s \leftarrow \bar{s}k \in \{0, \dots, 255\}^{32+32+32+32}$ 
4  $\mu' \leftarrow \text{NewHope-CPA-PKE.Decrypt}(c, sk)$ 
5  $K' || coin'' || d' \leftarrow \text{SHAKE-256}(96, \mu' || h) \in \{0, \dots, 255\}^{32+32+32}$ 
6 if  $c = \text{NewHope-CPA-PKE.Encrypt}(pk, \mu'; coin'')$  and  $d = d'$ 
7 then  $fail \leftarrow 0$  else  $fail \leftarrow 1$  end if
8  $K_0 \leftarrow K'$ 
9  $K_1 \leftarrow s$ 
10 return  $ss = (\text{SHAKE-256}(32, K_{fail} || \text{SHAKE-256}(32, c||d))$ 

```

Attack Model 2. We assume that Alice has done the CCA key generation. Eve sends messages m_E with a wrong coin. Alice will then reject any messages m_e because the verification is never passed. Eve, the attacker, is able to make a power analysis during the verification step of Alice's decapsulation algorithm.

With a power analysis, Eve can easily get the desired key mismatch oracle with a low number of traces.

A first idea would be to target the computation of $d = d'$ with differential power analysis. The following code corresponds to **NewHope-CCA-KEM** verification step where $a = (c, d)$, $b = (\text{NewHope-CPA-PKE.Encrypt}(pk, \mu'; coin''), d')$ and $len = 17/8 \cdot N + 32$.

```

* Name:          verify
* Description:   Compare two arrays for equality in constant time.
* Arguments:
* const unsigned char *a: pointer to first byte array
* const unsigned char *b: pointer to second byte array
* size_t len:    length of the byte arrays
* Returns 0 if the byte arrays are equal, 1 otherwise
*****/

int verify(const unsigned char *a, const unsigned char *b, size_t len)
{
    uint64_t r;
    size_t i;
    r = 0;

    for(i=0; i<len; i++)
        r |= a[i] ^ b[i];

    r = (-r) >> 63;
    return r;
}

```

This naive method actually works well in practice for an unprotected scheme because when $d = d'$, r is ored with 0 during $17/8 \cdot N$ iterations and when $d \neq d'$, r is ored with arbitrary values during $17/8 \cdot N$ iterations. With a single trace analysis, the equality $d = d'$ can be verified. One would argue reasonably that unprotected schemes are always vulnerable. The main aim of [21] is to propose a countermeasure to such an attack for a similar scheme which uses the CCA transform. It is an open problem, in this protected context, to extend this approach to a second order power analysis attack. A more realistic model relies on an invasive attack, this is what we present in the sequel.

Single Fault Attack. We consider inserting a fault during the computation of the verification step which cancels the CCA transform. The attack model becomes:

Attack Model 3. *We assume that Alice has done the CCA key generation. Eve, the attacker, is able to set the value r to 0 in the verification step of Alice's decapsulation algorithm.*

If Eve is able to set the value r to 0 anytime during the check $d = d'$, she can bypass the reencryption and the mismatch will appear only if $d \neq d'$. Then oracle \mathcal{O}_1 becomes accessible. Indeed, Eve can thus send any message m_E . Alice derives a wrong $coin'$ but the verification is skipped with the fault. If $d = d'$, Alice derives the shared key for initiating a communication. If $d \neq d'$, Alice will notice Eve that the key agreement failed. Eve will then deduce that the key is different. This vulnerability has been underlined in Sect. 3.6 of [21] for a similar scheme. But no countermeasure has been added to protect against this single fault attack, which can practically be induced by a laser. Countermeasures should have thus to be considered (see [18]), what may impact the efficiency of the verification.

6 Conclusion

The resilience of NIST post-quantum candidate algorithms in misuse situations is worth being investigated. It will help developers to propose implementations with countermeasures tightly designed to ensure the security in extreme contexts

(e.g. smart card, IoT) without decreasing too much the efficiency. In this paper, we describe an active attack against NewHope-CPA-KEM with (public, private) key pair reuse. This clearly confirms that if the designers' caveat against any private key reuse (e.g. temporary caching) is not strictly followed, this results in a practical, low complexity, key recovery attack. Our study indeed indicates that setting an upper limit of a few hundreds on the number of authorized key reuses would not be conservative enough, and already expose private keys to significant information leakages. While unprotected versions of CCA-KEM are extremely efficient, implementations of this scheme must be very carefully protected against any key mismatch oracle leakage if key pairs are potentially reused. As explained in this paper, this is particularly true for countermeasures against fault attacks. This might eventually come with a cost in terms of performance. This consideration may become even more important if one considers second order side channel or combined attacks, which could be a sequel of this work.

References

1. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange - a new hope. In: Holz, T., Savage, S. (eds.) 25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, 10–12 August 2016, pp. 327–343. USENIX Association (2016)
2. Bauer, A., Gilbert, H., Renault, G., Rossi, M.: Assessment of the Key-Reuse Resilience of NewHope (2019, to appear)
3. Bernstein, D.J., Groot Bruinderink, L., Lange, T., Panny, L.: HILA5 pindakaas: on the CCA security of lattice-based encryption with error correction. In: Joux, A., Nitaj, A., Rachidi, T. (eds.) AFRICACRYPT 2018. LNCS, vol. 10831, pp. 203–216. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89339-6_12
4. Bleichenbacher, D.: Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS #1. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 1–12. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0055716>
5. Bosma, W., Cannon, J., Playoust, C.: The Magma algebra system I. The user language. *J. Symbolic Comput.* **24**(3–4), 235–265 (1997). Computational algebra and number theory (London, 1993)
6. Braithwaite, M.: Experimenting with Post-quantum Cryptography. Posting on the Google Security Blog (2016)
7. Ding, J.: A Simple Provably Secure Key Exchange Scheme Based on the Learning with Errors Problem. Cryptology ePrint Archive, Report 2012/688 (2012). <https://eprint.iacr.org/2012/688>
8. Ding, J., Alsayigh, S., Saraswathy, R.V., Fluhrer, S.R., Lin, X.: Leakage of signal function with reused keys in RLWE key exchange. In: IEEE International Conference on Communications, ICC 2017, pp. 1–6. IEEE (2017)
9. Ding, J., Fluhrer, S., Saraswathy, R.V.: Complete attack on RLWE key exchange with reused keys, without signal leakage. In: Susilo, W., Yang, G. (eds.) ACISP 2018. LNCS, vol. 10946, pp. 467–486. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-93638-3_27
10. Fluhrer, S.: Cryptanalysis of Ring-LWE based Key Exchange with Key Share Reuse. Cryptology ePrint Archive, Report 2016/085 (2016). <https://eprint.iacr.org/2016/085>

11. Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 537–554. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48405-1_34
12. Guo, Q., Johansson, T., Stankovski, P.: A key recovery attack on MDPC with CCA security using decoding errors. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016, Part I. LNCS, vol. 10031, pp. 789–815. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53887-6_29
13. Hall, C., Goldberg, I., Schneier, B.: Reaction attacks against several public-key cryptosystem. In: Varadharajan, V., Mu, Y. (eds.) ICICS 1999. LNCS, vol. 1726, pp. 2–12. Springer, Heidelberg (1999). https://doi.org/10.1007/978-3-540-47942-0_2
14. Hoffstein, J., Silverman, J.H.: Protecting NTRU against chosen ciphertext and reaction attacks. Technical report 16, NTRU Cryptosystems Technical report (2000)
15. Hoffstein, J., Silverman, J.H.: Reaction attacks against the NTRU public key cryptosystem. Technical report 15, NTRU Cryptosystems Technical report (1999)
16. Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the fujisaki-okamoto transformation. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017. LNCS, vol. 10677, pp. 341–371. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70500-2_12
17. Howgrave-Graham, N., et al.: The impact of decryption failures on the security of NTRU encryption. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 226–246. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45146-4_14
18. Joye, M., Tunstall, M. (eds.): Fault Analysis in Cryptography. Information Security and Cryptography. Springer, Heidelberg (2012). <https://doi.org/10.1007/978-3-642-29656-7>
19. Kirkwood, D., Lackey, B., McVey, J., Motley, M., Solinas, J., Tuller, D.: Failure is not an option: standardization issues for post-quantum key agreement. In: NIST Workshop on Cybersecurity in a Post Quantum World (2015)
20. Menezes, A., Ustaoglu, B.: On reusing ephemeral keys in Diffie-Hellman key agreement protocols. IJACT **2**(2), 154–158 (2010)
21. Oder, T., Schneider, T., Pöppelmann, T., Güneysu, T.: Practical CCA2-secure and masked ring-LWE implementation. IACR Transactions on CHES (2016). <https://eprint.iacr.org/2016/1109>
22. Peikert, C.: Lattice cryptography for the internet. In: Mosca, M. (ed.) PQCrypto 2014. LNCS, vol. 8772, pp. 197–219. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11659-4_12
23. Pöppelmann, T., et al.: NewHope. Submission to Round 1 of NIST Post Quantum Cryptography Competition (2017)
24. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_23
25. Schwabe, P., Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: NewHope without Reconciliation. Cryptology ePrint Archive, Report 2016/1157 (2016). <https://eprint.iacr.org/2016/1157>