A provably masked implementation of BIKE Key Encapsulation Mechanism

Loïc Demange[1,2*] and Mélissa Rossi[3]

[1] Thales, Gennevilliers, France
[2] Inria, Paris, France `loic.demange@inria.fr`
[3] ANSSI, Paris, France `melissa.rossi@ssi.gouv.fr`

Abstract. BIKE is a post-quantum key encapsulation mechanism (KEM) selected for the 4th round of the NIST's standardization campaign. It relies on the hardness of the syndrome decoding problem for quasi-cyclic codes and on the indistinguishability of the public key from a random element, and provides the most competitive performance among round 4 candidates, which makes it relevant for future real-world use cases. Analyzing its side-channel resistance has been highly encouraged by the community and several works have already outlined various side-channel weaknesses and proposed ad-hoc countermeasures. However, in contrast to the well-documented research line on masking lattice-based algorithms, the possibility of generically protecting code-based algorithms by masking has only been marginally studied in a 2016 paper by Cong Chen et al. At this stage of the standardization campaign, it is important to assess the possibility of fully masking BIKE scheme and the resulting cost in terms of performances.

In this work, we provide the first high-order masked implementation of a code-based algorithm. We had to tackle many issues such as finding proper ways to handle large sparse polynomials, masking the key-generation algorithm or keeping the benefit of the bitslicing. In this paper, we present all the gadgets necessary to provide a fully masked implementation of BIKE, we discuss our different implementation choices and we propose a full proof of masking in the Ishai Sahai and Wagner (Crypto 2003) model.

More practically, we also provide an open C-code masked implementation of the key-generation, encapsulation and decapsulation algorithms with extensive benchmarks. While the obtained performance is slower than existing masked lattice-based algorithms, the scaling in the masking order is still encouraging and no Boolean to Arithmetic conversion has been used.

We hope that this work can be a starting point for future analysis and optimization.

Keywords: BIKE · PQC · Side-Channel countermeasure · Provable high-order masking · d-probing model

1 Introduction

In response to the potential quantum threat, the NIST has initiated a standardization campaign in 2017 for defining new post-quantum algorithms. Different families of mathematical problems have received a lot of attention. Particularly, lattices and error-correcting codes stood out as interesting building blocks for post-quantum schemes. Six years after the first round of the campaign, three lattice-based schemes have been selected as future NIST post-quantum standards. But in parallel, the standardization campaign continues for other key-encapsulation schemes like code-based ones, because it is important to be able to have diverse schemes based on other structures.

BIKE [ABB+22], a round 4 candidate for the NIST standardization process, is still under analysis by the research community. It relies on the hardness of the syndrome decoding problem for quasi-cyclic codes and on the indistinguishability of the public key from a random element. It is the most efficient and offers the lowest key sizes of all round 4 candidates. BIKE belongs to a line of research on QC-MDPC code-based schemes started by [MTSB13]. The advantage of QC-MDPC-based algorithms is the sparse structure of the underlying variables. Such codes allow for the use of iterative bit-flipping decoding algorithms (detailed later in the paper) as part of the decapsulation. The first implementations of QC-MDPC codes [vMG14, MOG15, vMHG16] were not constant

time and vulnerable to time attacks. In 2016, Chou proposed a portable constant-time C implementation [Cho16]. Some implementation proposals have been made over the years [DG17, GAB19, BOG20], and another constant-time C implementation was introduced in 2020 [DGK20], especially improving the decoding part, and which is claimed protected against timing and cache attacks. Cortex-M4 optimized implementations of BIKE have been introduced later in [CCK21]. Subsequently, optimizations have been proposed in [CGKT22].

Existing side-channel attacks   Timing vulnerabilities have been handled with priority in the previously stated implementations. However, the authors of [GHJ+21] have highlighted the possibility of using timing information of the constant weight word sampler in the decapsulation in order to apply [GJS16]'s reaction attack. Such a vulnerability has been thwarted by redesigning the word sampler in [Sen21].

On the power-consumption attacks side, several works have outlined various side-channel weaknesses and proposed ad-hoc countermeasures. Indeed, while BIKE's sparse and structured private keys are essential for providing good performances and compactness, this exact structure and redundancy can be exploited by side-channel attacks in order to lower down the difficulty of the underlying decoding problem. For instance, Chou's implementation has been targeted by a differential power analysis attack on the syndrome computation in [RHHM17]. Later, an improvement of the previous attack and a single-trace analysis exploiting leakage in the syndrome computation were provided in [SKC+19]. Very recently, [CARG23] introduced a new single-trace attack on the most recent implementation of BIKE. The authors use unsupervised clustering techniques on the trace during the cyclic shifts computation to recover some bits of the positions of the ones in the private key. Next, they combine such knowledge with classical information set decoding techniques to recover the full key.

Existing generic side-channel protections   The current implementations of code-based schemes are claimed to be protected against timing and cache attacks, but they are never fully masked, i.e. masked from key generation to decapsulation. Masking is known as the most deployed countermeasure against physical attackers and is widely applied in embedded systems. Masking consists in randomizing any secret-dependent intermediate variable. Each of these secret-dependent intermediate variables, say $\mathbf{x}$, is split into $d+1$ variables $(\mathbf{x}_i)_{0 \leq i \leq d}$ called "shares". The integer d is referred to as the masking order. In this paper, the only necessary type of masking is *Boolean masking*. In other words, a sensitive variable $\mathbf{x}$ is shared in $(\mathbf{x}_i)_{0 \leq i \leq d}$ such that

$$\mathbf{x} = \mathbf{x}_0 \oplus \cdots \oplus \mathbf{x}_d. \tag{1}$$

While $\mathbb{F}_2$-linear operations can straightforwardly be applied share-wise, non-linear operations are more complex and require additional randomness, as shown in [ISW03]. Proving the security of a masked design consists in showing that the joint distribution of any set of at most d intermediate variables is independent of the secrets. But, the bigger the algorithm is, the more dependencies to be considered in the proof. Fortunately, several works have defined intermediate security properties that simplify the security proofs [RP10, CPRR14, BBD+16]: one can focus on proving the properties on small parts of the algorithms, denoted gadgets, and it is possible to securely compose the pieces together.

Much effort has been performed on provably masking lattice-based primitives in the past five years and many challenges have been overcome. For example, [BBE+18] introduced a new security notion to justify unmasking certain intermediate steps. In [GR19], the authors proposed a masked implementation of the qTesla signature scheme [BAA+19]. In [KDV+22], a masked Fuijsaki-Okamoto transform is introduced for a fully masked Saber KEM implementation [DKR+20]. The NIST post-quantum finalists Crystals-Dilithium [LDK+22] and Crystals-Kyber [SAB+22] have also been masked in [ABC+22] and [BGR+21].

The picture is less abundant when it comes to code-based schemes. One explanation could come from the large sparse polynomials leading to potential prohibitive performances or the complex counter-based decoder. The authors of [KLRBG22a] propose a first-order masked inversion in multiplicative masking. Another recent work [KLRBG22b] presents a way to mask BIKE's key generation with a fixed weight polynomial sampling technique and arithmetic to Boolean conversions.

Our contribution   In this paper, we provide the first provable high-order masked implementation of a code-based algorithm. We detail every masked gadget that is necessary for masking BIKE's key generation and decapsulation. The proofs are given in the d-probing model. Let us detail some aspects of our design.

- **No mask conversion** Mask-conversion gadgets consist in modifying the underlying masking operation, e.g. going from $\oplus$ to an addition in $\mathbb{Z}_q$. Even if the unmasked functionality is the identity function, these

gadgets are known to be heavy in terms of computation time. Despite efficiency improvements since their introduction e.g. in [CGV14, Cor17, CGTV15], current secure mask conversion algorithms run in time at least $O(d^2)$. Contrary to lattices, BIKE is fundamentally relying on binary operations. While the authors of [KLRBG22b] have included mask conversion in their design, we believe that keeping only Boolean masking would be more natural and efficient. In this work, we give the first evidence that it is possible to completely mask BIKE without any mask conversion.

- **Sparse versus dense representation**. BIKE's intermediate variables are sparse polynomials with coefficients in $\mathbb{F}_2$. An important question arose rapidly when designing a masked BIKE: *Should we represent the masked polynomials in dense form or keep the sparse structure and mask the indices of the non-zero coefficient instead?* On the one hand, the number of non-zero coefficients is protected but the multiplication requires a masked Karatsuba-based multiplication algorithm. On the other hand, the number of non-zero coefficients is accessible by timing attacks but a lighter multiplication algorithm based on cyclic shifts is possible. The sparse representation intuitively seems lighter but some parts necessarily required the dense form for security. For completeness, we decided to analyze both following approaches:

  1. A fully-dense implementation where the polynomials are masked in dense form.

  2. A hybrid sparse-dense implementation where the polynomials are represented in sparse form whenever the number of non-zero coefficient is independent from any secret data.

Interestingly, our experiments showed that a fully-dense approach seems more relevant, especially for high orders. While (2) and (1) seem equivalent for one or two shares, (1) looks indeed more relevant for higher orders. This difference might shrink with more optimizations of the cyclic shift, as it will be discussed in the future work section.

- **Many new gadgets**. A lot of new gadgets needed to be introduced for masking BIKE. Although BIKE's bitslice addition technique turned out to operate well with Boolean masking, some other parts of the key generation were more challenging to mask. For example the Fisher-Yates sampling algorithm/technique and the polynomial inversion required many loops and subroutines. More generally, we provide in this paper all elementary gadgets that are necessary to mask BIKE even if their design did not pose any particular issue. We believe that they can be of independent interest for masking future code-based schemes.

We provide an open C-code implementation of the key-generation, encapsulation and decapsulation algorithms with detailed benchmarks. Although theoretically quadratic [ISW03], several post quantum masked designs can lead to an experimental scale in the masking order that tends to be exponential [BBE+18, Table 1]. The scaling we've obtained is very encouraging, as our experiments seem to indicate a quadratic scaling. We believe that it is even possible to further improve and optimize our code and maybe reach quasi-linearity in the masking order. We hope that this work can be a first building block towards masked code-based cryptography and could lead to future analysis and new optimization.

Organization of the paper   In Section 2, we introduce all the necessary background on masking, QC-MDPC codes and BIKE. In Section 3, we present our general masked construction along with its composition security proof. In Section 4, we detail the gadgets. For brevity, we only detail a few main gadgets and refer to Appendix A for the description of the remaining gadgets. Finally, in Section 5, we present our implementation and its benchmarks. We conclude with the future work in Section 6.

2   Preliminaries

2.1   Masking

A shared variable $(\mathbf{x}_i)_{0 \le i \le d}$ according to Eq. (1) will be denoted by $[\![\mathbf{x}]\!]$ for readability. Note that for a masking order d, there are $d+1$ shares.

Definition 1 (d-probing Security or ISW security [ISW03])**.** An algorithm is d-probing secure iff the joint distribution of any set of at most d internal intermediate values is independent of the secrets.

Even if d-probing security seems far from realistic side-channel protection, it is actually backed-up by theoretical model reductions that relate the d-probing security to side-channel security up to a certain level of noise [DDF14]. Moreover, [CJRR99] showed that the number of measurements required to mount a successful side-channel attack usually increases exponentially in the masking order.

In addition to Definition 1, other intermediate security properties were introduced to ease the security proofs [RP10, CPRR14, BBD+16]. The focus can be placed on proving these properties on small parts of the algorithms, denoted gadgets.

Definition 2 (Gadget). A gadget is a probabilistic algorithm that takes shared and unshared inputs values and returns shared and un-shared values.

These new security properties open the door for securely composing gadgets.

Definition 3 (Non interference [BBD+16]). A gadget is:

- d-non-interfering (d-NI) iff any set of at most d observations can be perfectly simulated from at most d shares of each input.

- d-strong non-interfering (d-SNI) iff any set of at most d observations whose d_{int} observations on the internal data and d_{out} observations on the outputs can be perfectly simulated from at most d_{int} shares of each input.

One can check that d-SNI implies d-NI, which itself implies d-probing security. Also note that any linear gadget for $\oplus$ is immediately d-NI.

2.2 Codes

In this paper, we will only introduce the relevant information for masking BIKE. Not many aspects of coding theory are needed for understanding our work.

Definition 4 (Binary linear codes). A binary linear code $\mathcal{C}$ of length n and dimension r is a r-dimensional vector subspace of $\mathbb{F}_2^n$. Is it possible to represent it in two equivalent ways:

- either using a generator matrix $\mathbf{G} \in \mathbb{F}_2^{r \times n}$ such that each row of $\mathbf{G}$ is an element of a basis of $\mathcal{C}$,

$$\mathcal{C} = \{m \cdot \mathbf{G}, m \in \mathbb{F}_2^r\}.$$

- or using a parity-check matrix $\mathbf{H} \in \mathbb{F}_2^{(n-r) \times n}$ such that for any $\mathbf{c} \in \mathcal{C}$,

$$\mathbf{c} \cdot \mathbf{H}^T = 0.$$

Definition 5 (Circulant matrix). An $r \times r$ matrix $\mathbf{A}$ is circulant if each row is a cyclic shift of the previous row. More precisely, $\mathbf{A}$ is of the form

$$\begin{pmatrix} a_0 & a_1 & \cdots & a_{r-1} \\ a_{r-1} & a_0 & \cdots & a_{r-2} \\ \vdots & & & \vdots \\ a_1 & a_2 & \dots & a_0 \end{pmatrix}.$$

We say that $\mathbf{A}$ is generated by the vector $(a_0, \cdots, a_{r-1})$.

Remark 1. It is possible to define an isomorphism between the ring of polynomials $\mathbb{F}_2[X]/(X^r - 1)$ and the set of circulant matrices of order r. To a vector $(a_0, \cdots, a_{r-1})$ generating a circulant matrix, one can associate the polynomial $\sum_{i=0}^{r-1} a_i X^i$. Multiplication and inversion can then be performed either with matrix multiplication or polynomial multiplication.

Definition 6 (Quasi-circulant matrix). A matrix is quasi-circulant if it is composed of circulant blocks.

For example, let

$$\mathbf{A} = \begin{pmatrix} a_1 & a_2 & a_3 \\ a_3 & a_1 & a_2 \\ a_2 & a_3 & a_1 \end{pmatrix} \text{ and } \mathbf{B} = \begin{pmatrix} b_1 & b_2 & b_3 \\ b_3 & b_1 & b_2 \\ b_2 & b_3 & b_1 \end{pmatrix}$$

be two circulant matrices. The matrix $\mathbf{C} = [\mathbf{A}|\mathbf{B}]$ defined as the concatenation of $\mathbf{A}$ and $\mathbf{B}$ is a quasi-circulant matrix.

Remark 2. Similarly to Remark 1, it is possible to represent quasi-circulant matrices as sets of polynomials.

Definition 7 (Quasi-cyclic code). A binary code $\mathcal{C}$ is quasi-cyclic iff it admits a quasi-circulant generating matrix.

Definition 8 (QC-MDPC code). Let n, r, w be integer parameters for length, dimension and minimum code weight. A $[n, r, w]$ QC-MDPC code $\mathcal{C}$ is a quasi-cyclic code that admits a parity-check matrix $\mathbf{H}$ such that $\mathbf{H}$ has a constant row weight $w = \mathcal{O}(\sqrt{n})$.

2.3   BIKE scheme

BIKE (*Bit Flipping Key Encapsulation*) [ABB+22] is a key encapsulation scheme based on QC-MDPC (Quasi-Cyclic Moderate Density Parity-Check) codes as introduced in Definition 8.

More precisely, let r and w be integer parameters. BIKE relies on $[2r, r, w]$ QC-MDPC codes. Its private key corresponds to the parity check matrix. The security of the scheme reduces to quasi-cyclic variants of hard problems from coding theory [Ale03, BMvT78]. We refer to [ABB+22] for more information about the security and design rationale.

BIKE's first building block is a public key encryption scheme (PKE) based on a variant of the Niederreiter framework [Nie86]. The plaintext is represented by the sparse vector (e_0, e_1), and the ciphertext by its syndrome. The decryption is performed with a decoding procedure that will be presented below in Section 2.4. Next, this PKE is converted into an IND-CCA KEM with the application of the Fujisaki-Okamoto transformation [HHK17]. This transformation requires several hash functions: $\mathbf{H}$, $\mathbf{L}$ and $\mathbf{K}$.

Let us detail the key generation (**KeyGen**), Encapsulation (**Encaps**) and Decapsulation (**Decaps**) algorithms in more details. In addition to the parameters r and w, let us define t and ℓ as integer parameters. We denote $\mathcal{R} = \mathbb{F}_2[X]/(X^r - 1)$ the underlying cyclic polynomial ring. Let us define

$$\begin{aligned} \mathcal{H}_w &= \{(h_0, h_1) \in \mathcal{R}^2 \mid |h_0| = |h_1| = w/2\}, \\ \mathcal{E}_t &= \{(e_0, e_1) \in \mathcal{R}^2 \mid |e_0| + |e_1| = t\}, \\ \mathcal{M} &= \{0, 1\}^\ell, \\ \mathcal{K} &= \{0, 1\}^\ell, \end{aligned}$$

as respectively the private key space, the error space, the message space and the shared key space. In the above, we denote by $|h|$ the Hamming weight of the polynomial h, i.e. the number of non-zero coefficients of h.

In the following, we denote $a \xleftarrow{\$} \mathcal{B}$ when a is sampled uniformly at random from $\mathcal{B}$, and $\leftarrow$ is an assigment of value.

Algorithm 1 Keygen

Ensure: $((h_0, h_1), \sigma) \in \mathcal{H}_w \times \mathcal{M}, h \in \mathcal{R}$

1: $(h_0, h_1) \xleftarrow{\$} \mathcal{H}_w$
2: $h \leftarrow h_1 h_0^{-1}$
3: $\sigma \xleftarrow{\$} \mathcal{M}$
4: **return** $((h_0, h_1, \sigma), h)$

Algorithm 2 Encaps

Require: $h \in \mathcal{R}$
Ensure: $K \in \mathcal{K}, c \in \mathcal{R} \times \mathcal{M}$

1: $m \xleftarrow{\$} \mathcal{M}$
2: $(e_0, e_1) \leftarrow \mathbf{H}(m)$
3: $c \leftarrow (e_0 + e_1 h, m \oplus \mathbf{L}(e_0, e_1))$
4: $K \leftarrow \mathbf{K}(m, c)$
5: **return** (K, c)

Algorithm 3 Decaps

Require: $(h_0, h_1, \sigma) \in \mathcal{H}_w \times \mathcal{M}$, $c = (c_0, c_1) \in \mathcal{R} \times \mathcal{M}$
Ensure: $K \in \mathcal{K}$
  1: $e' \leftarrow \mathsf{decoder}(c_0 h_0, h_0, h_1)$
  2: $m' \leftarrow c_1 \oplus \mathbf{L}(e')$
  3: **if** $e' = \mathbf{H}(m')$ **then**
  4:     $K \leftarrow \mathbf{K}(m', c)$
  5: **else**
  6:     $K \leftarrow \mathbf{K}(\sigma, c)$
  7: **end if**
  8: **return** K

Table 1: BIKE's proposed parameters [ABB$^+$22]

	Level 1	Level 3	Level 5
r	12323	24659	40973
w	142	206	274
t	134	199	264
ℓ	256	256	256

Parameter setting As defined in the specifications, the parameters should satisfy several constraints. The block length r should be a prime number, and 2 should be primitive modulo r. The parameter w should be such that $w = 2d \approx \sqrt{n}$ with d being odd. In addition, the error weight should be such that $t \approx \sqrt{n}$. We present the instantiated parameters in Table 1.

2.4 Decoding QC-MDPC codes

The choice of the decoder has a crucial impact on the security and the performances of the scheme. As QC-MDPC codes have sparse parity matrices, decoding techniques usually rely on Bit-Flipping techniques originally introduced in [Gal62] for low density parity-check matrices.

Technically, the Bit Flipping algorithm is presented in Algorithm 4 and works as follows: over several iterations, we compute the syndrome $c\mathbf{H}^T$ where c is the ciphertext and $\mathbf{H}^T$ is the transposed parity matrix of the code. Next, we count the number of unsatisfied parity-check equations for each position. If the counter for a position exceeds T, a pre-computed threshold (on the fly according to the weight of the syndrome), the position is flipped and the syndrome is recomputed. Let $\mathsf{syndrome}$ be the syndrome computation, $\mathsf{counter}$ the counter computation, and $\mathsf{threshold}$ the threshold computation function. We refer to [ABB$^+$22] for details.

Algorithm 4 Bit Flipping algorithm

Require: $\mathbf{H}^T$ the sparse parity matrix of a $[2r, r, w]$ MDPC code
     $c \in \mathbb{F}_2^n$ a noisy codeword
Ensure: A codeword c, $c\mathbf{H}^T = 0$
  1: $s \leftarrow \mathsf{syndrome}(c, \mathbf{H})$
  2: **while** $|s| \neq 0$ **do**
  3:     $T \leftarrow \mathsf{threshold}(|s|)$
  4:     **for** $j \in \{1, ..., n\}$ **do**
  5:         **if** $\mathsf{counter}(s, j, \mathbf{H}) \geq T$ **then**
  6:             $c_j \leftarrow c_j \oplus 1$
  7:         **end if**
  8:     **end for**
  9:     $s \leftarrow \mathsf{syndrome}(c, \mathbf{H})$
10: **end while**
11: **return** c

The authors of BIKE chose a refined Black-Gray-Flip (BGF) technique introduced in [DGK20]. This refinement classifies the error in a black zone or a gray zone using two different thresholds. Two additional iterations are performed to verify the choices made during the classification. The BFG decoding algorithm is presented in Algorithm 5. This decoder also has a fixed number of iterations (set at 5), to avoid timing attacks.

Algorithm 5 `Black-Gray-Flip` (BGF)

Parameters: $r, w, t, d = w/2, n = 2r$; Nbr_Iter, τ, `threshold` (see text for details)

Require: $s \in \mathbb{F}_2^r$, $H \in \mathbb{F}_2^{r \times n}$

1: $e \leftarrow 0^n$
2: **for** $i = 1, \ldots,$Nbr_Iter **do**
3:      $T \leftarrow$ `threshold`$(|s + e\mathbf{H}^T|, i)$
4:      $e, black, grey \leftarrow$ `BFIter`$(s + e\mathbf{H}^T, e, T, H)$
5:      **if** $i = 1$ **then**
6:          $e \leftarrow$ `BFMaskedIter`$(s + e\mathbf{H}^T, e, black, (d + 1)/2 + 1, H)$
7:          $e \leftarrow$ `BFMaskedIter`$(s + e\mathbf{H}^T, e, grey, (d + 1)/2 + 1, H)$
8:      **end if**
9: **end for**
10: **if** $s = e\mathbf{H}^T$ **then**
11:      **return** e
12: **else**
13:      **return** $\perp$
14: **end if**

15: **procedure** `BFIter`(s, e, T, H)
16: **for** $j = 0, \ldots, n - 1$ **do**
17:      **if** `ctr`$(H, s, j) \geq T$ **then**
18:          $e_j \leftarrow e_j \oplus 1$
19:          $black_j \leftarrow 1$
20:      **else if** `ctr`$(H, s, j) \geq T - \tau$ **then**
21:          $grey_j \leftarrow 1$
22:      **end if**
23: **end for**
24: **return** $e, black, grey$

25: **procedure** `BFMaskedIter`$(s, e, mask, T, H)$
26: **for** $j = 0, \ldots, n - 1$ **do**
27:      **if** `ctr`$(H, s, j) \geq T$ **then**
28:          $e_j \leftarrow e_j \oplus mask_j$
29:      **end if**
30: **end for**
31: **return** e

---

3 Masked BIKE

We present here the core contribution of this paper: a fully masked encapsulation, decapsulation and key generation for BIKE. While the encapsulation uses public data, most of it had to be masked anyway as part of the decapsulation process due to the IND-CCA transform. Thus, for a perfectly complete masked design, the masked encapsulation is also included in our code. The masked decapsulation is obviously the most important part as it is the primary target of side-channel attacks. A masked key generation can also be relevant to prevent single-trace key recovery attacks when the private key is generated. A masked encapsulation might be relevant in advanced attack models to prevent single-trace message-recovery attacks.

In this section, we present the salient ideas of our masking design. Details on some selected underlying gadgets will be presented later in Section 4. Some gadgets were already introduced in the literature but many new gadgets have been introduced to achieve our design. The complete list of gadgets is summed-up in Tables 2

Table 2: Security properties of the known gadgets.

Gadget	Security Property	Reference
$\mathsf{sec}_\&$	$d - \mathsf{SNI}$	[CGTV15, BBE$^+$18].
refresh	$d - \mathsf{SNI}$	[Cor14]
sec_+	$d - \mathsf{NI}$	[Cor14]

and 3.

3.1 Sparse and dense notation

BIKE's private key $\mathbf{H}$ is a sparse polynomial (see Remark 2). For masking such polynomials, both approaches are valid: either we represent in its dense form or we keep the sparse structure and mask the indices of the non-zero coefficients instead. Since the number of non-zero coefficient is a public parameter, two approaches are potentially valid. The sparse representation intuitively seems lighter but some part (such as error generation) will require the dense form for security reasons. For completeness, we analyze both approaches: (1) an implementation where $\mathbf{H}$ is masked in dense form and (2) a hybrid-sparse-dense implementation where both dense and sparse forms of $\mathbf{H}$ are stored.

The masked private key will then be denoted by $[\![\mathbf{h}_0^{\mathrm{sparse}}]\!]$, $[\![\mathbf{h}_1^{\mathrm{sparse}}]\!]$ when it is masked in sparse form (i.e. the indices of the non-zero coefficients are masked) and it will be denoted by $[\![\mathbf{h}_0^{\mathrm{dense}}]\!]$, $[\![\mathbf{h}_1^{\mathrm{dense}}]\!]$ when the full polynomial is masked. The same convention is applied for other intermediate variables that can be masked in dense or sparse form.

Let sparse_to_dense be an algorithm that converts the sparse notation into a dense notation by multiplying the sparse polynomial by a dense polynomial equal to 1. This procedure is straightforwardly d-NI.

3.2 Key generation

The masked key generation is introduced below in Algorithm 6. We use a masked version of the Fisher-Yates algorithm. It consists in drawing a vector of n random elements, where each position i contains a value between 0 and $n - i$. Since it is important to avoid any duplicates, we go through the array backwards and we replace the value by the index i in case of duplicates. Despite a bias in the distribution, this does not affect the security of the scheme as proved in [Sen21]. This will allow us to generate our private keys $\mathbf{h}_0$ and $\mathbf{h}_1$, to then compute the public key $\mathbf{h}$. Provided that all the gadgets enjoy the d-NI property, their sequential combination leads to a

---

Algorithm 6 Masked key generation

Ensure: $[\![\mathbf{h}_0^{\mathrm{sparse}}]\!] \in \mathbb{F}_2^{\frac{w}{2}}$, $[\![\mathbf{h}_1^{\mathrm{sparse}}]\!] \in \mathbb{F}_2^{\frac{w}{2}}$, $[\![\mathbf{h}_0^{\mathrm{dense}}]\!] \in \mathbb{F}_2^r$, $[\![\mathbf{h}_1^{\mathrm{dense}}]\!] \in \mathbb{F}_2^r$, $[\![\mathbf{h}]\!] \in \mathbb{F}_2^r$
1: $[\![\mathbf{h}_0^{\mathrm{sparse}}]\!] \leftarrow \mathsf{SecFisherYates}(\frac{w}{2}, r)$ ▷ Algorithm 17
2: $[\![\mathbf{h}_1^{\mathrm{sparse}}]\!] \leftarrow \mathsf{SecFisherYates}(\frac{w}{2}, r)$
3: $[\![\mathbf{h}_0^{\mathrm{dense}}]\!] \leftarrow \mathsf{sparse\_to\_dense}([\![\mathbf{h}_0^{\mathrm{sparse}}]\!])$ ▷ see Section 3.1
4: $[\![\mathbf{h}_1^{\mathrm{dense}}]\!] \leftarrow \mathsf{sparse\_to\_dense}([\![\mathbf{h}_1^{\mathrm{sparse}}]\!])$
5: $[\![\mathbf{h}_0^{-1}]\!] \leftarrow \mathsf{SecInversion}([\![\mathbf{h}_0^{\mathrm{dense}}]\!], r)$ ▷ Algorithm 18
6: $[\![\mathbf{h}]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathbf{h}_0^{-1}]\!], [\![\mathbf{h}_1^{\mathrm{dense}}]\!])$ ▷ Algorithm 13
7: **return** $\mathsf{sk} = \left([\![\mathbf{h}_0^{\mathrm{sparse}}]\!], [\![\mathbf{h}_1^{\mathrm{sparse}}]\!], [\![\mathbf{h}_0^{\mathrm{dense}}]\!], [\![\mathbf{h}_1^{\mathrm{dense}}]\!]\right)$, $\mathsf{pk} = [\![\mathbf{h}]\!]$

---

d-NI algorithm. Thus we have the following result.

Theorem 1. *The masked key generation algorithm is $d - \mathsf{NI}$.*

Table 3: Security properties of the introduced gadgets.

Non-Specific Gadgets

Gadget	d-NI Theorem	Function
SecAdder (Alg. 23 & 24)	Th. 16 & 17	Addition in $\mathbb{Z}$
$sec_=$ (Alg. 25)	Th. 18	Equality check
SecBitslice (Alg. 11 & 12)	Th. 6	Bitsliced addition
$SecMult_{partlymasked}$ (Alg. 26)	Th. 19	Multiplication between masked and unmasked
sec_{rand} (Alg. 27)	Th. 20	Modular random number
SecPolymul (Alg. 28)	Th. 21	Polynomial multiplication
SecKaratsuba (Alg.13)	Th. 7	Karatsuba multiplication
$sec_{\gg}$ (Alg. 29)	Th. 22	Cyclic shift
$SecMult_{sparsedense}$ (Alg. 14)	Th. 8	Multiplication between sparse and dense
sec_{hw} (Alg. 15)	Th. 9	Hamming weight computation
sec_{if} (Alg. 30)	Th. 23	Conditional if
sec_{max} (Alg. 31)	Th. 24	Maximum function
sec_{fill} (Alg. 32)	Th. 25	Matrix filling

BIKE-adapted Gadgets

Gadget	$d - $NI Theorem	Function
SecBGF (Alg. 9)	Th. 4	BGF Decoder (Alg. 5)
SecKeyGen (Alg. 6)	Th. 1	Key Generation (Alg. 1)
SecErrorGen (Alg. 7)	Th. 5	Generation of (e_0, e_1) in Decaps
SecGreyZone (Alg. 22)	Th. 15	Grey Zone technique ([DGK20])
SecFisherYates (Alg. 17)	Th. 11	Fisher-Yates generation of sparse polynomials ([Sen21])
SecInversion (Alg. 18)	Th. 12	Polynomial inversion
SecSyndrome (Alg. 16)	Th. 10	Syndrome computation
SecThreshold (Alg. 19)	Th. 13	BIKE's Threshold computation
SecCounter (Alg. 21)	Th. 14	BIKE's counter computation

3.3   Encapsulation

IND-CCA masked implementation   The IND-CCA security of the scheme is achieved thanks to the Fujisaki-Okamoto transformation. This transformation consists in XORing the seed used to generate the secret with the hashed secret. This will allow, during the decryption, to recover the seed and thus to check if the secret has been honestly generated. This transformation prevents active chosen ciphertext attack. In BIKE [ABB+22], the **K**, **L** and **H** hash functions (see Algorithm 3) are instantiated with SHAKE256 and SHA384. These functions have already been protected in the masked implementation of Saber (see [DKR+20] for more information about Saber) in [KDV+22]. This framework is easily adaptable for BIKE without major modification.

3.3.1   Error generation

The error generation algorithm is necessary for both encapsulation and decapsulation. Its masked version is introduced below in Algorithm 7. It consists in generating a masked error vector $[\![\mathbf{e}_0^{\mathrm{dense}}]\!], [\![\mathbf{e}_1^{\mathrm{dense}}]\!]$.
It uses two d-NI sub-gadgets:

- sec_+ corresponds to the logical addition of two integers [Cor14]. We introduce $\mathsf{sec}_{+\mathrm{partlymasked}}$ which is almost identical to sec_+ but where the first operation ($\mathsf{sec}_\&$ between the two masked parameters) has been modified to take an unmasked element (& between all parts of the masked value and the public one).
- $\mathsf{sec}_{\mathsf{if}}$ represents a conditional branch, it outputs either the first input or the second one depending on the Boolean value of the last input. It is detailed in Algorithm 30 in Appendix A.6.

The error cannot be represented in sparse representation, as the weights of $[\![\mathbf{e}_0^{\mathrm{dense}}]\!], [\![\mathbf{e}_1^{\mathrm{dense}}]\!]$ are not constant. This would leak sensitive information.

In this algorithm, the intermediate values are used only once within d-NI gadgets, the only exception being

---

Algorithm 7 Masked Error generation SecErrorGen

Require: $[\![s]\!] \in \mathbb{F}_2^{256}$ the seed for SecFisherYates
Ensure: $[\![\mathbf{e}_0^{\mathrm{dense}}]\!] \in \mathbb{F}_2^r$, $[\![\mathbf{e}_1^{\mathrm{dense}}]\!] \in \mathbb{F}_2^r$
1: $[\![\mathbf{e}^{\mathrm{sparse}}]\!] \leftarrow \mathsf{SecFisherYates}(t, 2 \times r)$ $\triangleright$ Algorithm 17
2: **for** $i \leftarrow 0$ to $t - 1$ **do**
3:    $[\![v]\!] \leftarrow \mathsf{sec}_{+\mathrm{partlymasked}}([\![\mathbf{e}_i^{\mathrm{sparse}}]\!], -r)$ $\triangleright$ see Section 3.3.1
4:    $[\![\mathbf{e}_i^{\mathrm{sparse}}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{e}_i^{\mathrm{sparse}}]\!])$
5:    $[\![s]\!] \leftarrow \mathsf{sec}_{\mathsf{if}}([\![\mathbf{e}_i^{\mathrm{sparse}}]\!], [\![v]\!], \mathsf{sign\_bit}([\![v]\!]))$ $\triangleright$ see Section 3.3.1
6:    $[\![\mathbf{t}]\!] \leftarrow \mathsf{sparse\_to\_dense}([\![s]\!])$ $\triangleright$ see Section 3.1, polynomial with only one coefficient
7:    $[\![\mathbf{e}_0^{\mathrm{dense}}]\!] \leftarrow [\![\mathbf{e}_0^{\mathrm{dense}}]\!] \oplus \mathsf{sec}_{\mathsf{if}}([\![\mathbf{t}]\!], \mathsf{vector\_zero\_masking}(), \mathsf{sign\_bit}([\![v]\!]))$ $\triangleright$ Coefficient-wise $\mathsf{sec}_{\mathsf{if}}$ and XOR
8:    $[\![\mathbf{e}_1^{\mathrm{dense}}]\!] \leftarrow [\![\mathbf{e}_1^{\mathrm{dense}}]\!] \oplus \mathsf{sec}_{\mathsf{if}}(\mathsf{vector\_zero\_masking}(), [\![\mathbf{t}]\!], \mathsf{sign\_bit}([\![v]\!]))$
9: **end for**
10: **return** $[\![\mathbf{e}_0^{\mathrm{dense}}]\!], [\![\mathbf{e}_1^{\mathrm{dense}}]\!]$

$[\![\mathbf{e}_i^{\mathrm{sparse}}]\!]$, which is refreshed ($d$-SNI) before its new use. We can therefore conclude with the following theorem.

Theorem 2. *The error generation algorithm is $d - \mathsf{NI}$.*

3.3.2   Encapsulation algorithm

All the functions used are $d - \mathsf{NI}$.
Since the only variable that has been reused is the seed **m** and the generated error **e**, we have to refresh them. We can conclude that the algorithm is itself $d - \mathsf{NI}$.

Theorem 3. *The encapsulation algorithm is $d - \mathsf{NI}$.*

3.4   Decapsulation

Decapsulation consists of first decoding the ciphertext and secondly checking that it is correct. While the most challenging masking work was on the decoding algorithm, we propose below a fully masked version of the decapsulation for completeness in Algorithm 10.

Algorithm 8 Encapsulation

Require: $[\![\mathbf{h}]\!] \in \mathbb{F}_2^r$
Ensure: $[\![\mathbf{c}]\!] \in \mathbb{F}_2^{r+32}$

1: $[\![\mathbf{m}]\!] \xleftarrow{\$} \mathbb{F}_2^{32}$
2: $[\![\mathbf{e}]\!] = ([\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!]) \leftarrow \mathsf{SecErrorGen}([\![\mathbf{m}]\!])$ ▷ Algorithm 7
3: $[\![\mathbf{m}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{m}]\!])$
4: $[\![\mathbf{c}]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathbf{e}_1]\!], [\![\mathbf{h}]\!])$
5: $[\![\mathbf{c}]\!] \leftarrow [\![\mathbf{c}]\!] \oplus [\![\mathbf{e}_0]\!]$ ▷ Coefficient-wise XOR
6: $[\![\mathbf{e}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{e}]\!])$
7: $[\![\mathbf{c}_r]\!] \leftarrow \mathsf{L}([\![\mathbf{e}]\!]) \oplus [\![\mathbf{m}]\!]$ ▷ see Section 3.3, Coefficient-wise XOR

3.4.1 BGF decoder

We now describe the most important part of the decapsulation: the masked BGF decoder. The unmasked version of the BGF decoder has been presented in Section 2.4. The masked version of Algorithm 5 is detailed in Algorithm 9. Recall that all the sub-gadgets are detailed in Section 4 and Appendix A (see Tables 2 and 3).

Where SecThreshold and SecSyndrome are fairly simple gadgets (based on additions, multiplications or shifts), SecCounter was quite challenging to mask as it relies on bitslicing. The SecGreyZone optimization allows for a much performant decoder but it also adds a layer of complexity in the decoding. This complexity is also transferred when masking is involved as several sensitive data are used inside the computations.

We denote by vector_zero_masking a subroutine that initializes a d sharing of an r-dimensional zero vector. Let $\mathbf{C}$ be a pair of matrices represented as table of dimension $2 \times r \times (\lfloor \frac{w}{2} \rfloor + 1)$, that can be decomposed into two matrices $\mathbf{C}_0$ and $\mathbf{C}_1$ of dimension $r \times \lfloor (\frac{w}{2}) + 1 \rfloor$. The notation $\mathbf{C}_{0,*,\lfloor \frac{w}{2} \rfloor}$ represents the entire row of height $\lfloor \frac{w}{2} \rfloor$.

Algorithm 9 BGF decoder

Require: $\mathsf{sk} = \left([\![\mathbf{h}_0^{\mathrm{dense}}]\!] \in \mathbb{F}_2^r, [\![\mathbf{h}_1^{\mathrm{dense}}]\!] \in \mathbb{F}_2^r, [\![\mathbf{h}_0^{\mathrm{sparse}}]\!] \in \mathbb{F}_2^{\frac{w}{2}}, [\![\mathbf{h}_1^{\mathrm{sparse}}]\!] \in \mathbb{F}_2^{\frac{w}{2}} \right), [\![\mathbf{c}_0]\!] \in \mathbb{F}_2^r$
Ensure: $[\![\mathbf{e}_0]\!] \in \mathbb{F}_2^r, [\![\mathbf{e}_1]\!] \in \mathbb{F}_2^r$ such that $(\mathbf{c}_0 + \mathbf{e}_0) \cdot \mathbf{h}_0 = 0$

1: $[\![\mathbf{e}_0]\!] \leftarrow \mathsf{vector\_zero\_masking}()$
2: $[\![\mathbf{e}_1]\!] \leftarrow \mathsf{vector\_zero\_masking}()$
3: $[\![\mathbf{s}]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathbf{c}_0]\!], [\![\mathbf{h}_0^{\mathrm{dense}}]\!])$ ▷ Algorithm 13
4: $[\![\mathbf{h}_0^{\mathrm{dense}}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{h}_0^{\mathrm{dense}}]\!])$
5: **for** $i \leftarrow 0$ **to** Nbr_Iter -1 **do**
6:      $[\![\mathbf{s1}]\!] \leftarrow \mathsf{SecSyndrome}([\![\mathbf{h}_0^{\mathrm{dense}}]\!], [\![\mathbf{h}_1^{\mathrm{dense}}]\!], [\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!], [\![\mathbf{s}]\!])$ ▷ Algorithm 16
7:      $[\![T]\!] \leftarrow \mathsf{SecThreshold}([\![\mathbf{s1}]\!])$ ▷ Algorithm 19
8:      $[\![\mathbf{s1}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{s1}]\!])$
9:      $[\![\mathbf{C}]\!] \leftarrow \mathsf{SecCounter}([\![\mathbf{s1}]\!], [\![T]\!], [\![\mathbf{h}_0^{\mathrm{sparse}}]\!], [\![\mathbf{h}_1^{\mathrm{sparse}}]\!])$ ▷ Algorithm 21
10:     $[\![\mathbf{e}_0]\!] \leftarrow \mathsf{refresh}([\![\mathbf{e}_0]\!]); [\![\mathbf{e}_1]\!] \leftarrow \mathsf{refresh}([\![\mathbf{e}_1]\!])$
11:                  ▷ $[\![\mathbf{C}_{0,*,\lfloor \frac{w}{2} \rfloor}]\!]$ and $[\![\mathbf{C}_{1,*,\lfloor \frac{w}{2} \rfloor}]\!]$ are the sign bit of the counters minus the threshold
12:     $[\![\mathbf{e}_0]\!] \leftarrow \neg(([\![\mathbf{e}_0]\!]) \oplus ([\![\mathbf{C}_{0,*,\lfloor \frac{w}{2} \rfloor}]\!]))$ ▷ Coefficient-wise XOR
13:     $[\![\mathbf{e}_1]\!] \leftarrow \neg(([\![\mathbf{e}_1]\!]) \oplus ([\![\mathbf{C}_{1,*,\lfloor \frac{w}{2} \rfloor}]\!]))$ ▷ Coefficient-wise XOR
14:     **if** $i = 0$ **then**
15:        $[\![\mathbf{s}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{s}]\!])$
16:        $[\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!] \leftarrow \mathsf{SecGreyZone}([\![\mathbf{C}]\!], [\![\mathbf{h}_0^{\mathrm{dense}}]\!], [\![\mathbf{h}_1^{\mathrm{dense}}]\!], [\![\mathbf{h}_0^{\mathrm{sparse}}]\!], [\![\mathbf{h}_1^{\mathrm{sparse}}]\!], [\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!], [\![\mathbf{s}]\!])$ ▷ Algorithm 22
17:     **end if**
18:     $[\![\mathsf{sk}]\!] \leftarrow \mathsf{refresh}([\![\mathsf{sk}]\!])$
19:     $[\![\mathbf{s}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{s}]\!])$
20: **end for**
21: **return** $([\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!])$

Theorem 4. *The BGF decoder algorithm is* $d - \mathsf{NI}$.

Proof. We represent the whole decoding algorithm in Figs. 1 and 2. To avoid complex graphs, the content of

an iteration for $i \neq 0$ can be proved separately (if $i \neq 0$, there is no application of the SecGreyZone algorithm, in Lines 14 to 16).
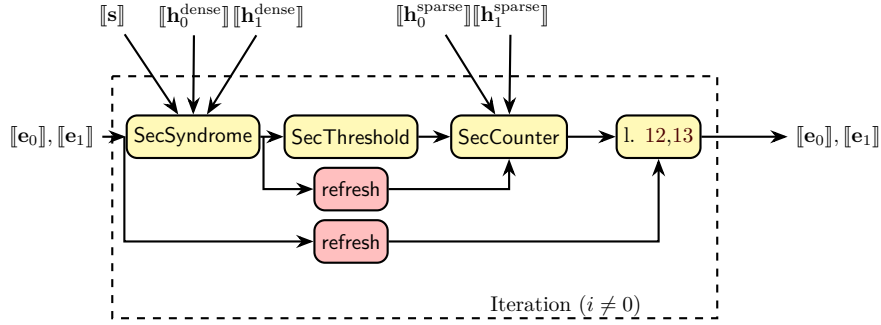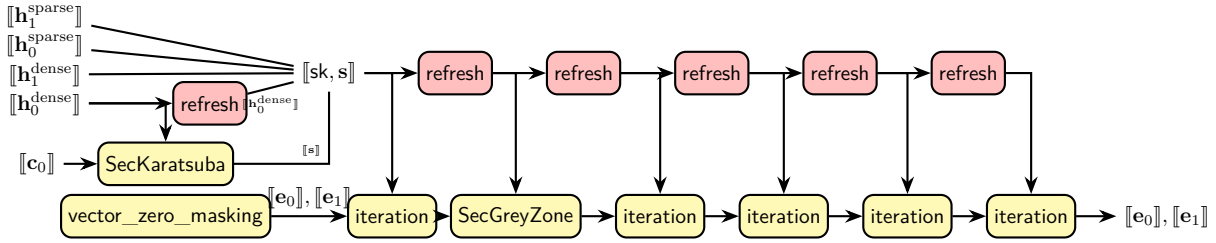

Figure 1: Structure of an iteration

Figure 2: Structure of the BGF decoder

Let us first look at one iteration with $i \neq 0$. Let us assume that it is a gadget with inputs $[\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!], [\![\mathsf{sk}]\!]$ and $[\![\mathbf{s}]\!]$. And we assume that this iteration's output is a modified version of $[\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!]$. Let us assume that an attacker has access to $\delta \leq d$ observations on this sub-gadget. Thus, we want to prove that all these δ observations can be perfectly simulated with at most δ shares of $[\![\mathsf{sk}]\!], [\![\mathbf{s}]\!], [\![\mathbf{e}_0]\!]$ and $[\![\mathbf{e}_1]\!]$. To fix notations, let us consider the following distribution of the attacker's δ observations:

- δ_6 on Lines 12 and 13,
- δ_5 during the SecCounter computation,
- δ_4 during the SecThreshold computation,
- δ_3 during the SecSyndrome computation,
- δ_2 when $[\![\mathbf{s}1]\!]$ si refreshed,
- δ_1 when $[\![\mathbf{e}_0]\!]$ and $[\![\mathbf{e}_1]\!]$ are refreshed.

By definition of the d-probing model, we have $\sum_{j=1}^{6} \delta_i \leq \delta \leq d$.

Since Lines 12 and 13 are $\mathbb{F}_2$-linear operations performed share by share, this computation verifies the d-NI property. In addition, all the gadgets are either $d - $ NI or $d - $ SNI as specified in Table 3. The proofs will be provided later in the paper. Finally, all the observations performed during this iteration can be perfectly simulated with at most $\sum_{j=1}^{6} \delta_i$ shares of $[\![\mathbf{e}_0]\!]$, the same amount for $[\![\mathbf{e}_1]\!]$, $\delta_6 + \delta_5$ shares of $\mathbf{h}_0^{\mathrm{sparse}}$, the same for $\mathbf{h}_0^{\mathrm{sparse}}$, $\sum_{j=2}^{6} \delta_i$ shares of $\mathbf{h}_0^{\mathrm{dense}}$ and finally the same for $\mathbf{h}_1^{\mathrm{dense}}$.

In the end, we have proved that all the probes can be perfectly simulated with at most $\delta \leq d$ shares of $[\![\mathsf{sk}]\!], [\![\mathbf{s}]\!], [\![\mathbf{e}_0]\!]$ and $[\![\mathbf{e}_1]\!]$.

Now let us analyze the complete construction in Fig. 2. The same reasoning applies. Let us assume that an attacker has access to $\delta \leq d$ observations on this algorithm. We consider the following distribution of the attacker's δ observations:

- $\delta_{\mathrm{iter,i}}$ on each $i - th$ iteration,
- $\delta_{\mathsf{SecGreyZone}}$ on the SecGreyZone computation,
- $\delta_{\mathrm{ref,i}}$ on the $i - th$ refresh of the secret key and the syndrome,

- $\delta_{\mathsf{vector\_zero\_masking}}$ on the vector_zero_masking computation,
- $\delta_{\mathsf{SecKaratsuba}}$ on the computation of the syndrome,
- δ_{ref} on the very first refresh.

By definition, $\sum_{i=0}^{Nbr\_Iter-1} (\delta_{\mathrm{iter,i}} + \delta_{\mathrm{ref,i}}) + \delta_{\mathsf{SecGreyZone}} + \delta_{\mathsf{vector\_zero\_masking}} + \delta_{\mathsf{SecKaratsuba}} + \delta_{\mathrm{ref}} \leq \delta \leq d$.
All the gadgets are proved d-NI and the refresh gadgets are d-SNI. All the probes performed after the first iteration (including the grey zone, the key refresh and the other following iterations), can be perfectly simulated with at most $\sum_{i=0}^{Nbr\_Iter-1} (\delta_{\mathrm{iter,i}} + \delta_{\mathrm{ref,i}}) + \delta_{\mathsf{SecGreyZone}}$ shares of $[\![\mathsf{sk}]\!]$, $[\![\mathbf{s}]\!]$, $[\![\mathbf{e}_0]\!]$ and $[\![\mathbf{e}_1]\!]$. Next, we use the probing security of the refresh, SecKaratsuba and vector_zero_masking. All the probes performed during the full decoding algorithm can be perfectly simulated with at most $\sum_{i=0}^{Nbr\_Iter-1} (\delta_{\mathrm{iter,i}} + \delta_{\mathrm{ref,i}}) + \delta_{\mathsf{SecGreyZone}} + \delta_{\mathsf{SecKaratsuba}} + \delta_{\mathrm{ref}}$ shares of $[\![\mathbf{c}_0]\!]$, the same for $[\![\mathbf{h}_0^{\mathrm{dense}}]\!]$ and $\sum_{i=0}^{Nbr\_Iter-1} (\delta_{\mathrm{iter,i}} + \delta_{\mathrm{ref,i}}) + \delta_{\mathsf{SecGreyZone}}$ for the rest of the secret key. All these numbers are smaller than to $\delta \leq d$ which concludes the proof.

$\square$

3.4.2 Decapsulation algorithm

For the needs of the decapsulation algorithm, we will introduce subvector fonction, an algorithm which returns the subvector starting and ending with the bounds given as parameters.

Algorithm 10 Decapsulation

Require: $[\![\mathbf{c}]\!] \in \mathbb{F}_2^{r+32}$, $[\![\vec{\sigma}]\!] \in \mathbb{F}_2^{32}$
Ensure: $[\![\mathbf{k}]\!] \in \mathbb{F}_2^{32}$
 1: $[\![\mathbf{e}']\!] \leftarrow \mathsf{SecBGF}([\![\mathbf{h}_0^{\mathrm{dense}}]\!], [\![\mathbf{h}_1^{\mathrm{dense}}]\!], [\![\mathbf{h}_0^{\mathrm{sparse}}]\!], [\![\mathbf{h}_1^{\mathrm{sparse}}]\!], [\![\mathbf{c}]\!])$ $\triangleright$ Algorithm 9
 2: $[\![\mathbf{m}']\!] \leftarrow \mathsf{L}([\![\mathbf{e}']\!])$ $\triangleright$ see Section 3.3
 3: $[\![\mathbf{m}']\!] \leftarrow [\![\mathbf{m}']\!] \oplus \mathsf{subvector}([\![\mathbf{c}]\!], r, r+32)$ $\triangleright$ see Section 3.4.2, coefficient-wise XOR
 4: $([\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!]) \leftarrow \mathsf{SecErrorGen}([\![\mathbf{m}']\!])$ $\triangleright$ Algorithm 7
 5: $[\![\mathbf{m}']\!] \leftarrow \mathsf{refresh}([\![\mathbf{m}']\!])$
 6: $[\![v]\!] \leftarrow 1$ $\triangleright$ Masked value of 1
 7: **for** $i \leftarrow 0$ to 1 **do**
 8:      **for** $j \leftarrow 0$ to $r-1$ **do**
 9:          $[\![t]\!] \leftarrow \mathsf{sec}_=([\![\mathbf{e}_{i,j}]\!], [\![\mathbf{e}'_{i,j}]\!])$ $\triangleright$ see Appendix A.2
10:          $[\![t]\!]_0 \leftarrow [\![t]\!]_0 \oplus 1$
11:          $[\![v]\!] \leftarrow \mathsf{sec}_\&([\![v]\!], [\![t]\!])$
12:      **end for**
13: **end for**
14: $[\![\mathbf{t}]\!] \leftarrow \mathsf{K}([\![\mathbf{m}']\!], [\![\mathbf{c}]\!])$ $\triangleright$ Section 3.4.2
15: $[\![\mathbf{t}1]\!] \leftarrow \mathsf{K}([\![\vec{\sigma}]\!], [\![\mathbf{c}]\!])$
16: $[\![\mathbf{k}]\!] \leftarrow \mathsf{sec}_{\mathsf{if}}([\![\mathbf{t}]\!], [\![\mathbf{t}1]\!], [\![v]\!])$ $\triangleright$ Coefficient-wise $\mathsf{sec}_{\mathsf{if}}$

Algorithm 10 uses $d-\mathsf{NI}$ gadgets, and the only variable that is used twice without modification is $\mathbf{m}'$. However, the dependancy loop is broken by the $d-\mathsf{SNI}$ refresh. Thus, we introduce the following theorem.

Theorem 5. *The decapsulation algorithm is $d-\mathsf{NI}$.*

4 Details on selected gadgets

In this section, we provide some details about selected gadgets.

4.1 Bitslicing

Bitslicing was introduced in [Cho16] for the QcBits implementation, with many similarities to BIKE. These techniques allow computations to be performed very efficiently and in constant time by focusing on the binary representation. In Algorithms 11 and 12, we present two versions of this BitSlice procedure depending on the type of the input. Both versions will be used in our implementation.

In Algorithms 11 and 12, we denote by SecHalf_Adder the procedure that computes the addition in $\mathbb{Z}$ of the inputs while outputing the carry as a second output. The SecAdder performs the same operation but is

given an extra carry. These simple gadgets are detailed and proved d-NI in Appendix A.1 for completeness. We also denote by zero_masking an initialization of a d-sharing of zero.

Algorithm 11 SecHalf_Bitslice

Require: $[\![\mathbf{X} := (\mathbf{X}_0, \cdots, \mathbf{X}_\ell)]\!] \in \mathbb{Z}_2^{k \times \ell}$, $[\![\mathbf{y}]\!] \in \mathbb{Z}_2^\ell$
Ensure: $[\![\mathbf{X}]\!] \in \mathbb{Z}_2^{\ell \times k}$ the result of the bitsliced addition between $[\![\mathbf{X}]\!]$ and $\mathbf{y}$
1: **for** $i := 0$ to $\ell - 1$ **do**
2:    $[\![r]\!] := [\![\mathbf{y}_i]\!]$
3:    **for** $j := 0$ to $k - 1$ **do**
4:       $([\![\mathbf{X}_{ij}]\!], [\![r]\!]) \leftarrow$ SecHalf_Adder$([\![\mathbf{X}_{ij}]\!], [\![r]\!])$
5:    **end for**
6: **end for**
7: **return** $[\![\mathbf{X}]\!]$

Algorithm 12 SecBitslice

Require: $[\![\mathbf{X} := (\mathbf{X}_0, \cdots, \mathbf{X}_k)]\!] \in \mathbb{Z}_2^{\ell \times k}$, $[\![\mathbf{Y} := (\mathbf{Y}_0, \cdots, \mathbf{Y}_k)]\!] \in \mathbb{Z}_2^{\ell \times k}$
Ensure: $[\![\mathbf{X}]\!] \in \mathbb{Z}_2^{\ell \times k}$ the result of the bitsliced addition between $[\![\mathbf{X}]\!]$ and $[\![\mathbf{Y}]\!]$
1: **for** $i := 0$ to $\ell - 1$ **do**
2:    $[\![r]\!] \leftarrow$ zero_masking()
3:    **for** $j := 0$ to $k - 1$ **do**
4:       $([\![\mathbf{X}_{ij}]\!], [\![r]\!]) \leftarrow$ SecAdder$([\![\mathbf{X}_{ij}]\!], [\![\mathbf{Y}_{ij}]\!], [\![r]\!])$
5:    **end for**
6: **end for**
7: **return** $[\![\mathbf{X}]\!]$

Since both SecHalf_Adder and SecAdder are d-NI and all loop iterations use different or updated variables, their sequential combination leads to a d-NI algorithm. Hence the following theorem.

Theorem 6. *The* SecHalf_Bitslice *and* SecBitslice *algorithms are* $d - $NI.

4.2 Multiplications

Several multiplication algorithms are necessary for masking BIKE. Indeed, as opposed to many other masked designs, the multiplication often takes two masked inputs instead of only one. In addition, the underlying $\mathbb{F}_2$ structure makes NTT-based multiplications irrelevant in BIKE's context. Thus, one valid solution is to fully mask the classical Karatsuba algorithm, as presented below. We denote by SecPolymul the naive schoolbook polynomial multiplication (detailed in Algorithm 28 in Appendix A.3 for completeness). Let B be a parameter denoting the recursion depth. It is fixed experimentally to allow performance optimization. In our experiments, we have fixed $B = 64$. We also set a parameter $s \in \mathbb{N}$ as a power of two corresponding to the size of the inputs. Let split be a subroutine that splits the $s/2$ high order and $s/2$ low order bits into two variables.

Theorem 7. *The Karatsuba algorithm is* $d - $NI *for any power of two* s *and any bound* $B \leq s$.

Proof. Let us prove this theorem by induction on the parameter s. If $s \leq B$, the $d - $NI property is directly inherited from the $d - $NI property of SecPolymul (Theorem 21 in Appendix A.3). Let us assume that the Karatsuba algorithm is $d - $NI for $s > B$ and let us sketch a proof that is it $d - $NI for the next power of two: $2 \cdot s$. The algorithm first computes $[\![\mathbf{z}_1]\!], [\![\mathbf{z}_2]\!]$ with $d - $NI gadgets. Then, the dependencies are broken by the d-SNI refresh before computing $[\![\mathbf{z}_3]\!]$. Finally, the recombination of $[\![\mathbf{z}_1]\!], [\![\mathbf{z}_2]\!]$ and $[\![\mathbf{z}_3]\!]$ uses only coefficient-wise $\mathbb{F}_2$-linear operations. Thus, Karatsuba algorithm is $d - $NI for $2 \cdot s$ which concludes the proof. $\square$

Remark 3 (Generalization to arbitrary s). Note that it is possible to generalize Karatsuba for multiplying two polynomials of any degree s. This generalization can be obtained with an extra padding before the multiplication and a modulo application afterwards. Since the size of polynomials and padding is public and the padding will itself be masked, this does not raise any security concerns. In this paper, we use the same notation "SecKaratsuba" even when the multiplication is applied in the context of polynomials.

Algorithm 13 Karatsuba multiplication on vectors

Require: $[\![\mathbf{p}1]\!] \in \mathbb{F}_2^s$, $[\![\mathbf{p}2]\!] \in \mathbb{F}_2^s$
Ensure: $[\![\mathbf{z}]\!] = [\![\mathbf{p}1]\!] \cdot [\![\mathbf{p}2]\!] \in \mathbb{F}_2^{2s}$
 1: **if** $s = B$ **then**
 2:     **return** SecPolymul($[\![\mathbf{p}1]\!]$, $[\![\mathbf{p}2]\!]$)$\triangleright$ Naive polynomial multiplication, see Algorithm 28 in Appendix A.3
 3: **end if**
 4: $([\![\mathsf{left}1]\!], [\![\mathsf{right}1]\!]) \leftarrow \mathsf{split}([\![\mathbf{p}1]\!])$ $\triangleright$ Splitting the $s/2$ high order and $s/2$ low order bits
 5: $([\![\mathsf{left}2]\!], [\![\mathsf{right}2]\!]) \leftarrow \mathsf{split}([\![\mathbf{p}2]\!])$ $\triangleright$ Splitting the $s/2$ high order and $s/2$ low order bits
 6: $[\![\mathbf{z}1]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathsf{right}1]\!], [\![\mathsf{right}2]\!])$
 7: $[\![\mathbf{z}2]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathsf{left}1]\!], [\![\mathsf{left}2]\!])$
 8: $[\![\mathsf{left}1]\!] \leftarrow \mathsf{refresh}([\![\mathsf{left}1]\!])$
 9: $[\![\mathsf{right}1]\!] \leftarrow \mathsf{refresh}([\![\mathsf{right}1]\!])$
10: $[\![\mathsf{left}2]\!] \leftarrow \mathsf{refresh}([\![\mathsf{left}2]\!])$
11: $[\![\mathsf{right}2]\!] \leftarrow \mathsf{refresh}([\![\mathsf{right}2]\!])$
12: $[\![\mathbf{t}1]\!] \leftarrow [\![\mathsf{left}1]\!] \oplus [\![\mathsf{right}1]\!]$ $\triangleright$ Coefficient-wise XOR
13: $[\![\mathbf{t}2]\!] \leftarrow [\![\mathsf{left}2]\!] \oplus [\![\mathsf{right}2]\!]$ $\triangleright$ Coefficient-wise XOR
14: $[\![\mathbf{z}3]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathbf{t}1]\!], [\![\mathbf{t}2]\!])$
15: **return** $[\![\mathbf{z}]\!] \leftarrow [\![\mathbf{z}1]\!] \oplus ([\![\mathbf{z}2]\!] \ll s/4) \oplus ([\![\mathbf{z}3]\!] \ll s/2)$ $\triangleright$ Coefficient-wise

---

In parallel, we also introduce a multiplication algorithm that takes only one masked input, the other input being a public value, as this algorithm is also necessary for our design. We denote it $\mathsf{SecMult}_{\mathsf{partlymasked}}$ and its design is detailed in Appendix A.3. It is directly inspired from the Montgomery ladder technique.

Leveraging sparse polynomials   In BIKE, it is often possible to leverage the fact that some masked polynomials are stored in sparse notation. We then introduce an extra gadget that takes one masked dense input and one masked sparse input. The multiplication technique uses a cyclic shift, denoted $\mathsf{sec}_{\gg}$. The idea is to shift a masked dense polynomial by a masked value. It is described and proved in Appendix A.5.

Algorithm 14 Sparse-dense multiplication ($\mathsf{SecMult}_{\mathsf{sparsedense}}$)

Require: $[\![\mathbf{x}]\!] \in \mathbb{F}_2^n$, $[\![\mathbf{y}]\!] \in \mathbb{F}_2^c$
Ensure: $[\![\mathbf{z}]\!] = [\![\mathbf{x}]\!] \cdot [\![\mathbf{y}]\!] \in \mathbb{F}_2^n$
 1: **for** $i \leftarrow 0$ to $c - 1$ **do**
 2:     $[\![\mathbf{t}]\!] \leftarrow \mathsf{sec}_{\gg}([\![\mathbf{x}]\!], [\![\mathbf{y}_i]\!])$ $\triangleright$ See Algorithm 29 in Appendix A.5
 3:     $[\![\mathbf{x}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{x}]\!])$
 4:     $[\![\mathbf{z}]\!] \leftarrow [\![\mathbf{z}]\!] \oplus [\![\mathbf{t}]\!]$ $\triangleright$ Coefficient-wise XOR
 5: **end for**
 6: **return** $[\![\mathbf{z}]\!]$

Theorem 8. *The* $\mathsf{SecMult}_{\mathsf{sparsedense}}$ *algorithm is* $d - \mathsf{NI}$.

Proof. Since the gadgets $\mathsf{sec}_{\gg}$ and $\oplus$ are d-NI. And even if $\mathbf{x}$ is reused in each loop, $\mathbf{x}$ is refreshed (d-SNI). □

4.3   Hamming weight

We introduce a masked Hamming weight computation. It has been optimized and involves the masked bitslice algorithm presented in Algorithm 12. Similarly to Karatsuba, we denote by right and left the cut in length of the matrix. For example, if $[\![\mathbf{T}]\!] \in \mathbb{F}_2^{l \times k}$, $\mathsf{right}([\![\mathbf{T}]\!])$ and $\mathsf{left}([\![\mathbf{T}]\!]) \in \mathbb{F}_2^{\frac{l}{2} \times k}$.
$[\![\mathbf{T}]\!]$ is a matrix that starts with one row, and will gain one more row per loop turn (call to bitslice). So we initialize $[\![\mathbf{T}_0]\!]$ as a vector, then at each iteration, $[\![\mathbf{T}]\!]$ will gain a row.

Algorithm 15 Hamming weight ($\mathsf{sec_{hw}}$)

Require: $[\![\mathbf{x}]\!] \in \mathbb{F}_2^n$
Ensure: $[\![y]\!] \in \mathbb{F}_2$ the hamming weight of $[\![\mathbf{x}]\!]$
1: $[\![\mathbf{T}_0]\!] \leftarrow [\![\mathbf{x}]\!]$ ▷ We initialize the first line of the $[\![\mathbf{T}]\!]$ matrix with $[\![\mathbf{x}]\!]$ vector
2: $j \leftarrow 1$
3: **for** $i \leftarrow \frac{n}{2}$ to 1 step $\frac{i}{2}$ **do**
4:      $[\![\mathbf{T}]\!] \leftarrow \mathsf{SecBitslice}(\mathsf{left}([\![\mathbf{T}]\!]), \mathsf{right}([\![\mathbf{T}]\!]))$ ▷ Cut in length
5:      $j \leftarrow j + 1$
6: **end for**
7: $[\![y]\!] \leftarrow \mathsf{zero\_masking}()$
8: **for** $i \leftarrow 0$ to $j - 1$ **do**
9:      $[\![y]\!] \leftarrow [\![y]\!] \oplus ([\![\mathbf{T}_{0,i}]\!] \ll i)$
10: **end for**
11: **return** $[\![y]\!]$

Theorem 9. *The hamming weight algorithm is $d - \mathsf{NI}$.*

Proof. Since as $\mathsf{SecBitslice}$ has been proved d-NI in Theorem 6 and all loops use updated variables, their composition leads to a d-NI algorithm. $\qquad\square$

In this part we will introduce the main gadgets necessary for the realization of masked BIKE.

4.4   Computing the syndrome

Algorithm 16 compute_syndrome

Require: $[\![\mathbf{h}_0^{\mathrm{dense}}]\!] \in \mathbb{F}_2^R$, $[\![\mathbf{h}_1^{\mathrm{dense}}]\!] \in \mathbb{F}_2^R$, $[\![\mathbf{e}_0]\!] \in \mathbb{F}_2^R$, $[\![\mathbf{e}_1]\!] \in \mathbb{F}_2^R$, $[\![\mathbf{s}]\!] = c_0 h_0$
Ensure: $[\![\mathbf{s}1]\!] = c_0 h_0 + e_0 h_0 + e_1 h_1 \in \mathbb{F}_2^r$
1: $[\![\mathbf{s}2]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathbf{e}_0]\!], [\![\mathbf{h}_0^{\mathrm{dense}}]\!])$
2: $[\![\mathbf{s}3]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathbf{e}_1]\!], [\![\mathbf{h}_1^{\mathrm{dense}}]\!])$
3: $[\![\mathbf{s}1]\!] \leftarrow [\![\mathbf{s}]\!] \oplus [\![\mathbf{s}2]\!] \oplus [\![\mathbf{s}3]\!]$ ▷ Coefficient-wise XOR
4: **return** $[\![\mathbf{s}1]\!]$

Since different variables are used in each of the function calls (all $d - \mathsf{NI}$), we get the following theorem.

Theorem 10. *The syndrome computing algorithm is $d - \mathsf{NI}$.*

4.5   Generation of random polynomials

The generation of sparse polynomials is performed using the Fisher-Yates technique. It was already introduced in Section 3.2. This procedure can be masked as presented in Algorithm 17. It uses $\mathsf{sec_{rand}}$, presented in Appendix A.4.

Algorithm 17 Fisher-Yates (SecFisherYates)

Require: $s \in \mathbb{N}$, $n \in \mathbb{N}$
Ensure: $[\![\mathbf{r}]\!] \in \mathbb{Z}_n^s$ a randomly generated vector without repeated values
 1: **for** $i \leftarrow s - 1$ to 0 **do**
 2:    $[\![\mathbf{r}_i]\!] \leftarrow \mathsf{sec_{rand}}(n - i)$
 3:    Initialize $[\![i]\!]$ as a Boolean sharing of i
 4:    $[\![\mathbf{r}_i]\!] \leftarrow \mathsf{sec_{+partlymasked}}([\![\mathbf{r}_i]\!], i)$
 5:    **for** $j \leftarrow i + 1$ to $s - 1$ **do**
 6:       $[\![\mathbf{r}_j]\!] \leftarrow \mathsf{refresh}([\![\mathbf{r}_j]\!])$
 7:       $[\![b]\!] \leftarrow \mathsf{sec_=}([\![\mathbf{r}_i]\!], [\![\mathbf{r}_j]\!])$
 8:       $[\![\mathbf{r}_i]\!] \leftarrow \mathsf{refresh}([\![\mathbf{r}_i]\!])$
 9:       $[\![i]\!] \leftarrow \mathsf{refresh}([\![i]\!])$
10:       $[\![\mathbf{r}_i]\!] \leftarrow \mathsf{sec_{if}}([\![i]\!], [\![\mathbf{r}_i]\!], [\![b]\!]))$
11:    **end for**
12: **end for**
13: **return** $[\![\mathbf{r}]\!]$

Theorem 11. *The Fisher-Yates algorithm is $d - \mathsf{NI}$.*

Proof. The Fisher-Yates algorithm involves many dependency loops. Indeed, each random $[\![\mathbf{r}_i]\!]$ is compared to all the previously derived ones. However, each value is refreshed before being used. Thus, the loop in lines 6 to 10 can be seen itself as a d-SNI gadget outputting $[\![\mathbf{r}_i]\!]$. Besides, the operations in lines 2 to 4 are d-NI. Hence, the outer loop can be seen as a sequential combination of NI gadgets and a d-SNI gadget for lines 6 to 10. In consequence, the algorithm is $d - \mathsf{NI}$. $\qquad\square$

4.6 Masked polynomial inversion

A masked polynomial inversion is needed for inverting $\mathbf{h}_0$ inside the key generation. The masked polynomial inversion is presented in Algorithm 18.

We note $\mathsf{sec_{pow}}$ a d-NI gadget allowing to raise a polynomial to the given (known) power. Since we only perform elevations of powers of 2, it boils down to permutations as the underlying ring is $\mathbb{F}_2$.

Algorithm 18 SecInversion

Require: $[\![\mathbf{x}]\!] \in \mathbb{F}_2^n$
Ensure: $[\![\mathbf{y}]\!] = [\![\mathbf{x}]\!]^{-1} \in \mathbb{F}_2^n$
 1: $[\![\mathbf{f}]\!] \leftarrow [\![\mathbf{x}]\!]$
 2: $[\![\mathbf{y}]\!] \leftarrow [\![\mathbf{x}]\!]$
 3: $[\![\mathbf{y}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{y}]\!])$
 4: **for** $i \leftarrow 0$ to $|n| - 1$ **do**
 5:    $[\![\mathbf{g}]\!] \leftarrow \mathsf{sec_{pow}}([\![\mathbf{f}]\!], 2^{2^i})$
 6:    $[\![\mathbf{f}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{f}]\!])$
 7:    $[\![\mathbf{f}]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathbf{f}]\!], [\![\mathbf{g}]\!])$
 8:    **if** the $(i+1)^{\text{th}}$ bit of $n - 2$ is 1 **then**
 9:       $[\![\mathbf{t}]\!] \leftarrow \mathsf{sec_{pow}}([\![\mathbf{f}]\!], 2^{(n-2)\ (\text{mod } 2^{(i+1)})})$
10:       $[\![\mathbf{y}]\!] \leftarrow \mathsf{SecKaratsuba}([\![\mathbf{y}]\!], [\![\mathbf{t}]\!])$
11:    **end if**
12: **end for**
13: $[\![\mathbf{y}]\!] \leftarrow \mathsf{sec_{pow}}([\![\mathbf{y}]\!], 2)$
14: **return** $[\![\mathbf{y}]\!]$

Theorem 12. *The masked inversion algorithm is $d - \mathsf{NI}$.*

Proof. The first iteration of the algorithm is presented in Fig. 3. One can graphically conclude that each iteration is d-NI as all the observations can be simulated with at most d shares of $([\![\mathbf{f}]\!], [\![\mathbf{y}]\!])$. Thus, the full loop
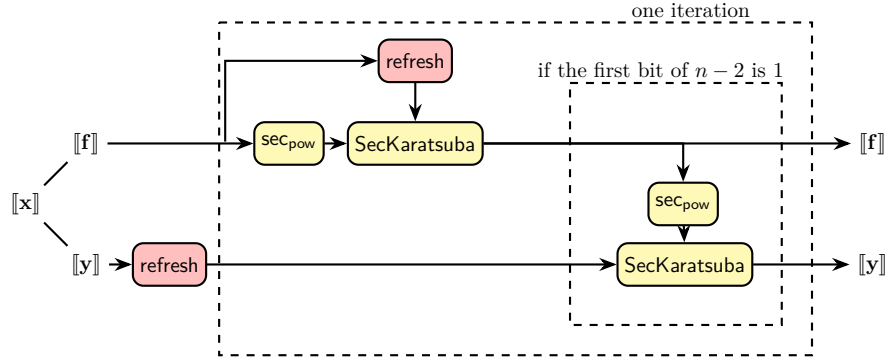
Figure 3: Sub-structure of the polynomial inversion algorithm

is d-NI. In addition, the final operation is d-NI. And, since both $[\![\mathbf{f}]\!]$ and $[\![\mathbf{y}]\!]$ are initialized with the same input $[\![\mathbf{x}]\!]$, one of them should be refreshed to end up with a full d-NI gadget.

$\square$

4.7 Threshold and counters

The decoder needs the computation of a threshold and counters, as presented in Algorithms 19 and 20. The threshold is an integer value that needs to be recomputed several times during decoding. Initially, the calculation of the threshold is done with floats, which is a concern for the masking. We have therefore reduced this to simple operations on integers such as threshold is equal to $\mathsf{max}(\lfloor \frac{T_0 \cdot S + T_1}{2^{T_2}} \rfloor, T_3)$.

The procedure to mask it involves gadgets previously introduced apart from $\mathsf{sec}_{\mathsf{max}}$, a gadget that, given two masked values, computes the greatest. The $\mathsf{sec}_{\mathsf{max}}$ gadget is detailed in Algorithm 31 in Appendix A.7.

Algorithm 19 SecThreshold

Require: $[\![\mathbf{s}]\!] \in \mathbb{F}_2^r$
Ensure: $[\![T]\!] \in \mathbb{Z}_r$ the threshold calculated from the syndrome
1: $[\![S]\!] \leftarrow \mathsf{sec}_{\mathsf{hw}}([\![\mathbf{s}]\!])$ $\triangleright$ Algorithm 15
2: $[\![T]\!] \leftarrow \mathsf{sec}_{\mathsf{T}}([\![S]\!])$ $\triangleright$ Algorithm 20
3: **return** $[\![T]\!]$

Algorithm 20 T computing ($\mathsf{sec}_{\mathsf{T}}$)

Require: $[\![S]\!] \in \mathbb{Z}_r$, T_0, T_1, T_2, T_3 fixed parameters of the scheme
Ensure: $[\![T]\!] = \mathsf{max}(\lfloor \frac{T_0 \cdot S + T_1}{2^{T_2}} \rfloor, T_3) \in \mathbb{Z}_r$
1: $[\![t]\!] \leftarrow \mathsf{SecMult}_{\mathrm{partlymasked}}([\![S]\!], T_0)$ $\triangleright$ Algorithm 26
2: $[\![T]\!] \leftarrow \mathsf{sec}_{+\mathrm{partlymasked}}([\![t]\!], T_1)$
3: $[\![T]\!] \leftarrow [\![T]\!] \gg T_2$
4: $[\![T]\!] \leftarrow \mathsf{sec}_{\mathsf{max}}([\![T]\!], [\![T_3]\!])$ $\triangleright$ Algorithm 31
5: **return** $[\![T]\!]$

Since we perform a sequence of operations that are d-NI themselves, we can establish the following theorem.

Theorem 13. *The computation of the threshold is* $d - \mathsf{NI}$.

During decoding, it is necessary to compute the number of unsatisfied parity check equations. We present in Algorithm 21 a masked version of this routine. Let denote by $[\![\mathbf{C}]\!] \in \mathbb{F}_2^{2 \times r \times (\lfloor \frac{w}{2} \rfloor + 1)}$ the matrix containing the binary representations of the counters of each coefficient. We manipulate this matrix as two double dimensional matrices, $[\![\mathbf{C}_0]\!]$ and $[\![\mathbf{C}_1]\!]$. Let $\mathsf{matrix\_zero\_masking}$ be the initialization of a d-sharing of a 2-dimensional zero matrix. This algorithm uses a gadget that consists in filling a matrix with a value. This technical gadget does not present any difficulties and is detailed in Algorithm 32 in Appendix A.8.

Theorem 14. *The counter computing algorithm is* $d - \mathsf{NI}$.

Proof. The procedure in lines 7 to 9 of Algorithm 21 is depicted in Fig. 4. One can see that all the loops are broken with a d-SNI refresh gadget. Thus lines 7 to 9 can be seen as a d-NI gadget.

The rest of the algorithm is a sequence of d-NI gadgets ($\mathsf{SecMult}_{\mathrm{partlymasked}}$, $\mathsf{sec}_{\mathsf{fill}}$, SecBitslice), thus the full algorith is $d - \mathsf{NI}$. $\square$

Algorithm 21 Counter computing (SecCounter)

Require: $[\![\mathbf{s}]\!] \in \mathbb{F}_2^r$, $[\![T]\!] \in \mathbb{Z}_r$, $[\![\mathbf{h}_0^{\mathrm{sparse}}]\!] \in \mathbb{F}_2^{\frac{w}{2}}$, $[\![\mathbf{h}_1^{\mathrm{sparse}}]\!] \in \mathbb{F}_2^{\frac{w}{2}}$

Ensure: $[\![\mathbf{C}]\!] \in \mathbb{F}_2^{2 \times r \times (|\frac{w}{2}|+1)}$ two matrices containing the binary representations of the counters for each coefficient

1: $[\![-T]\!] \leftarrow \mathsf{SecMult}_{\mathrm{partlymasked}}([\![T]\!], -1)$ ▷ Algorithm 26
2: $[\![\mathbf{C}_0]\!] \leftarrow \mathsf{matrix\_zero\_masking}()$
3: $[\![\mathbf{C}_1]\!] \leftarrow \mathsf{matrix\_zero\_masking}()$
4: $[\![\mathbf{P}]\!] \leftarrow [[\![\mathbf{h}_0^{\mathrm{sparse}}]\!], [\![\mathbf{h}_1^{\mathrm{sparse}}]\!]]$
5: **for** $i \leftarrow 0$ to 1 **do**
6:      **for** $j \leftarrow 0$ to $\frac{w}{2} - 1$ **do**
7:          $[\![\mathbf{s}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{s}]\!])$
8:          $[\![\mathbf{z}]\!] \leftarrow \mathsf{sec}_{\gg}([\![\mathbf{s}]\!], [\![\mathbf{P}_{i,j}]\!])$ ▷ Algorithm 29
9:          $[\![\mathbf{C}_i]\!] \leftarrow \mathsf{SecHalf\_Bitslice}([\![\mathbf{C}_i]\!], [\![\mathbf{z}]\!])$ ▷ Algorithm 11
10:      **end for**
11: **end for**
12: $[\![\mathbf{T}_0]\!] \leftarrow \mathsf{sec\_fill}([\![-T]\!])$ ▷ Algorithm 32
13: $[\![\mathbf{T}_1]\!] \leftarrow \mathsf{sec\_fill}(\mathsf{refresh}([\![-T]\!]))$ ▷ Algorithm 32
14: $[\![\mathbf{C}_0]\!] \leftarrow \mathsf{SecBitslice}([\![\mathbf{C}_0]\!], [\![\mathbf{T}_0]\!])$ ▷ Algorithm 12
15: $[\![\mathbf{C}_1]\!] \leftarrow \mathsf{SecBitslice}([\![\mathbf{C}_1]\!], [\![\mathbf{T}_1]\!])$ ▷ Algorithm 12
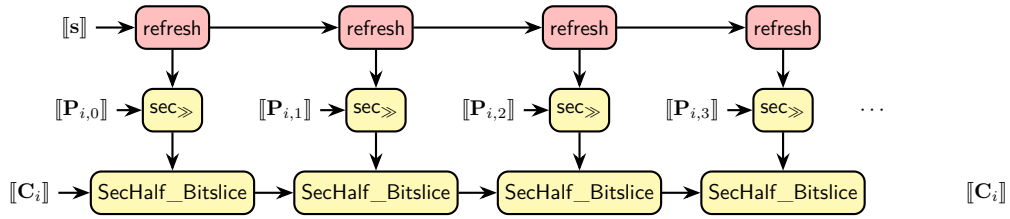16: **return** $[\![\mathbf{C}]\!] = [[\![\mathbf{C}_0]\!], [\![\mathbf{C}_1]\!]]$

Figure 4: Sub-structure of the counter algorithm

4.8 Grey Zone

The grey zone is an additional iteration of the decoder that is only realized at the first loop of the decoder. We will carry out the same operations as the classic decoder, but with an additional step with another threshold in order to detect more false positions and to be able to catch the possible errors of some ambiguous positions.

It takes as input the black zone $\mathbf{T}_0$, which is the matrix containing the counters minus the threshold. With, we can calculate the grey zone $\mathbf{T}_1$, which contains the counters minus the threshold plus τ.

Theorem 15. *The grey zone algorithm is $d - \mathsf{NI}$.*

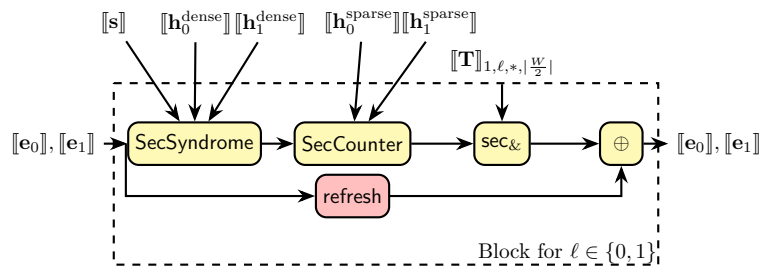*Proof.* We define a particular gadget called "block" for lines 11 to 17.

Figure 5: Structure of one block

Algorithm 22 SecGreyZone

Require: $\mathsf{sk} = \left([\![\mathbf{h}_0^{\mathrm{dense}}]\!] \in \mathbb{F}_2^r, [\![\mathbf{h}_1^{\mathrm{dense}}]\!] \in \mathbb{F}_2^r, [\![\mathbf{h}_0^{\mathrm{sparse}}]\!] \in \mathbb{F}_2^{\frac{w}{2}}, [\![\mathbf{h}_1^{\mathrm{sparse}}]\!] \in \mathbb{F}_2^{\frac{w}{2}} \right)$, $[\![\mathbf{T}_0]\!] \in \mathbb{F}_2^{2 \times r \times (\lfloor \frac{w}{2} \rfloor + 1)}$, $[\![\mathbf{e}_0]\!] \in \mathbb{F}_2^r$,
    $[\![\mathbf{e}_1]\!] \in \mathbb{F}_2^r$, $[\![\mathbf{s}]\!] \in \mathbb{F}_2^r$
Ensure: $[\![\mathbf{e}_0]\!] \in \mathbb{F}_2^r$, $[\![\mathbf{e}_1]\!] \in \mathbb{F}_2^r$

  1: Initialize $[\![\tau]\!]$ as a Boolean sharing of 3   ▷ 3 is a fixed parameter
  2: $[\![\mathbf{V}]\!] \leftarrow \mathsf{sec}_{\mathsf{fill}}([\![\tau]\!])$ ▷ Algorithm 32
  3: $[\![\mathbf{T}_{1,0}]\!] \leftarrow \mathsf{SecBitslice}([\![\mathbf{T}_{0,0}]\!], [\![\mathbf{V}]\!])$ ▷ Algorithm 12
  4: $[\![\mathbf{V}]\!] \leftarrow \mathsf{refresh}([\![\mathbf{V}]\!])$
  5: $[\![\mathbf{T}_{1,1}]\!] \leftarrow \mathsf{SecBitslice}([\![\mathbf{T}_{0,1}]\!], [\![\mathbf{V}]\!])$ ▷ Algorithm 12
  6: $[\![\mathbf{T}_0]\!] \leftarrow \mathsf{refresh}([\![\mathbf{T}_0]\!])$
  7: $[\![\mathbf{T}_{1,0,*,\lfloor \frac{w}{2} \rfloor}]\!] \leftarrow [\![\mathbf{T}_{0,0,*,\lfloor \frac{w}{2} \rfloor}]\!] \oplus [\![\mathbf{T}_{1,0,*,\lfloor \frac{w}{2} \rfloor}]\!]$ ▷ Coefficient-wise XOR
  8: $[\![\mathbf{T}_{1,1,*,\lfloor \frac{w}{2} \rfloor}]\!] \leftarrow [\![\mathbf{T}_{0,1,*,\lfloor \frac{w}{2} \rfloor}]\!] \oplus [\![\mathbf{T}_{1,1,*,\lfloor \frac{w}{2} \rfloor}]\!]$ ▷ Coefficient-wise XOR
  9: **for** $l \leftarrow 0$ to 1 **do**
 10:     $[\![\mathsf{sk}]\!] \leftarrow \mathsf{refresh}([\![\mathsf{sk}]\!])$
 11:     $[\![\mathbf{s1}]\!] \leftarrow \mathsf{SecSyndrome}([\![\mathbf{h}_0^{\mathrm{dense}}]\!], [\![\mathbf{h}_1^{\mathrm{dense}}]\!], [\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!], [\![\mathbf{s}]\!])$ ▷ Algorithm 16
 12:     $[\![\mathbf{C}]\!] \leftarrow \mathsf{SecCounter}([\![\mathbf{s1}]\!], \frac{\frac{w}{2}+1}{2}, [\![\mathbf{h}_0^{\mathrm{sparse}}]\!], [\![\mathbf{h}_1^{\mathrm{sparse}}]\!])$ ▷ Algorithm 21
 13:     $[\![\mathbf{v}_0]\!] \leftarrow \mathsf{sec}_\&(\neg[\![\mathbf{C}_{0,*,\lfloor \frac{w}{2} \rfloor}]\!], [\![\mathbf{T}_{l,0,*,\lfloor \frac{w}{2} \rfloor}]\!])$ ▷ Coefficient-wise $\mathsf{sec}_\&$
 14:     $[\![\mathbf{v}_1]\!] \leftarrow \mathsf{sec}_\&(\neg[\![\mathbf{C}_{1,*,\lfloor \frac{w}{2} \rfloor}]\!], [\![\mathbf{T}_{l,1,*,\lfloor \frac{w}{2} \rfloor}]\!])$ ▷ Coefficient-wise $\mathsf{sec}_\&$
 15:     $[\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!] \leftarrow \mathsf{refresh}([\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!])$
 16:     $[\![\mathbf{e}_0]\!] \leftarrow [\![\mathbf{e}_0]\!] \oplus [\![\mathbf{v}_0]\!]$ ▷ Coefficient-wise XOR
 17:     $[\![\mathbf{e}_1]\!] \leftarrow [\![\mathbf{e}_1]\!] \oplus [\![\mathbf{v}_1]\!]$ ▷ Coefficient-wise XOR
 18: **end for**
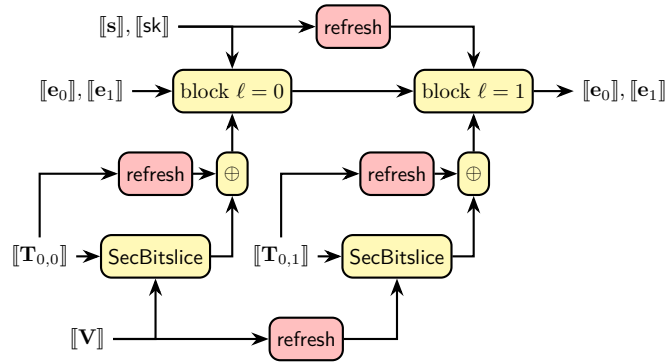 19: **return** $[\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!]$

Figure 6: Structure of the grey zone gadget

The overall details of the dependencies are presented in Figs. 5 and 6. As illustrated in Fig. 5, the block gadget is d-NI. Indeed, the only dependency loop is broken by a d-SNI refresh algorithm. Let us consider the full algorithm. Let us assume that an attacker has access to $\delta \leq d$ observations on this gadget. Then, we want to prove that all these δ observations can be perfectly simulated with at most δ shares of $[\![\mathsf{sk}]\!], [\![\mathbf{T}_0]\!], [\![\mathbf{e}_0]\!], [\![\mathbf{e}_1]\!], [\![\mathbf{s}]\!]$ and $[\![\mathbf{V}]\!]$ (note that the last one can be omitted as it is derived from a public parameter). To fix notations, let us consider the following distribution of the attacker's δ observations:

- δ_1 during the bitslice of Line 3
- δ_2 during the refreshing of $\mathbf{V}$
- δ_3 during the bitslice of Line 5
- δ_4 during the refresh of $\mathbf{T}_0$ (splitted in two sub-gadgets in the figure)
- δ_5 during the $\oplus$ in Line 8,
- δ_6 during the $\oplus$ in Line 7,
- δ_7 during the refreshing of sk and $\mathbf{s}$,
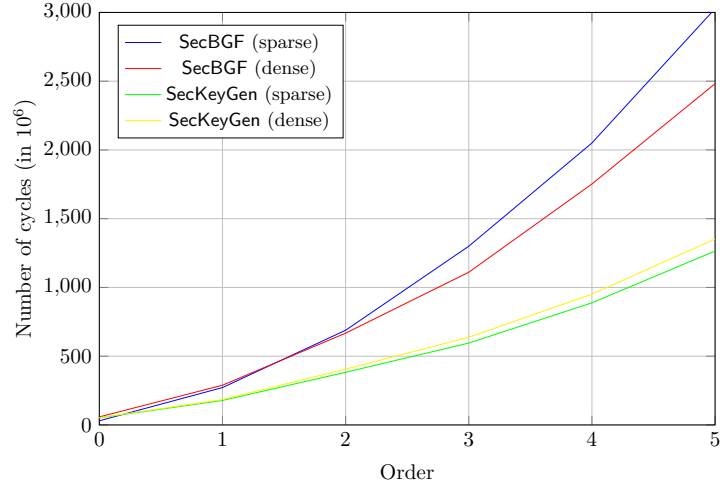- δ_8 in the block with $\ell = 0$,

Figure 7: Sparse vs dense, SecBGF and SecKeyGen

- δ_9 in the block with $\ell = 1$

By definition of the d-probing model, we have $\sum_{j=1}^{9} \delta_i \leq \delta \leq d$. All the gadgets are proved d-NI and the refresh gadgets are d-SNI. We skip the progressive part of the proof and directly claim that all the observations that are made during the execution of the gadget can be perfectly simulated with

- $\delta_2 + \delta_9 + \delta_8 + \delta_5 + \delta_1 + \delta_7$ shares of $[\![\mathbf{V}]\!]$
- $\delta_9 + \delta_8 + \delta_7$ shares of $[\![\mathsf{sk}]\!]$ and $[\![\mathbf{s}]\!]$ (each)
- $\delta_9 + \delta_8 + \delta_5 + \delta_4 + \delta_1$ shares of $\mathbf{T}_{0,0}$,
- $\delta_9 + \delta_6 + \delta_3 + \delta_4$ shares of $\mathbf{T}_{0,1}$,
- $\delta_9 + \delta_8$ shares of $[\![\mathbf{e}_0]\!]$ and $[\![\mathbf{e}_1]\!]$ (each).

This can be verified with the help of the figure. All these numbers of shares are smaller to $\delta \leq d$ which concludes the proof. □

5 Performance and experiments

5.1 Implementation

All the gadgets introduced in this paper have been implemented in large and complete C-code. Side-channel attacks are highly dependent on the chip on which the algorithm is executed and it is true that assembly codes are always the best pratical solution. However, C-code seems the best option to provide a multi-platform proof of concept. This code could be reused for future analysis and optimizations. The full code will be publicly available for code-checking and reproducibility. You can find it on Github[1].

Sparse vs dense representation Since most of the computations are polynomial operations performed on sparse objects, let us recall that we had two available options: the fully-dense implementation and the hybrid-sparse-dense one. In the first case, we see the polynomials as dense (with a conversion of the keys during the SecKeyGen) and we use Karatsuba for the majority of the calculations. In the other case, since we can represent a number of polynomials in sparse representation, we use SecMult$_{\mathsf{sparsedense}}$ as much as possible.
As presented in Fig. 7, our benchmarks show that while both approaches seem equivalent for one or two shares, a fully dense approach is indeed more relevant for higher orders.

One can conclude from our work that for the moment (except with potentially upcoming new optimizations), the dense representation seems more relevant. We will therefore keep the dense representation for the rest of the benchmarks, as it scales better when the order exceeds or equals 2.

---

[1] https://github.com/loicdemange/masked_BIKE_code

Table 4: Scaling benchmarks on particular gadgets, i7-4710MQ 2.5Ghz gcc 12.2.0 -03, NIST Level 1, median results on 200 executions

	Order 0	1	2	3	4	5
SecBGF(Alg. 9) (sparse)	1	×9	×22.8	×43	×67.9	×100.2
SecBGF(Alg. 9) (dense)	1	×4.9	×11.4	×19	×30	×42.4
SecKeyGen (Alg. 6) (dense)	1	×3.5	×7.7	×12.1	×18	×25.6
SecErrorGen (Alg. 7)	1	×8.6	×22.3	×40.8	×66	×96.2
SecGreyZone(Alg. 22) (sparse)	1	×9	×23.9	×41.9	×67.4	×98.3
SecGreyZone(Alg. 22) (dense)	1	×4.8	×11.2	×18.8	×29.2	×42
SecFisherYates(Alg. 17)	1	×9.5	×18.5	×29.7	×45.7	×66
SecInversion(Alg. 18)	1	×3.5	×7.2	×11	×16.4	×23.6
SecSyndrome(Alg. 16) (sparse)	1	×8	×21.3	×42.3	×63.7	×93
SecSyndrome(Alg. 16) (dense)	1	×3.3	×7.3	×10.8	×15.7	×22.2
SecThreshold(Alg. 19)	1	×8.6	×12.9	×19.1	×30.5	×43.1
SecCounter(Alg. 21)	1	×9.4	×23.1	×42.1	×67.3	×97
SecKaratsuba(Alg. 13)	1	×3.4	×7.4	×11.1	×16.7	×23.4
SecMult$_{sparsedense}$(Alg. 14)	1	×8.2	×21.3	×40	×64.8	×95.6

5.2   Detailed benchmarks

The code was benchmarked on an i7-4710MQ running at 2.5Ghz, 8GB of RAM, and compiling with gcc 12.2.0 -O3 flag. The given performances are obtained for NIST security level 1 ($r = 12323$). Identical experiments can provide data for the other security levels. Multiple benchmarks were performed and the results are listed in Table 4. We can notice that the performance of the gadgets depends on the performance of the multiplicative gadgets.

Bottlenecks   The sparse-dense multiplication and the basic Karatsuba multiplication seem to be the bottleneck of our implementation. An optimization of these gadgets could lead to big improvements of the complete scheme's performance. One idea to improve these gadgets could be to optimize the last recursion calculation of Karatsuba. In the unmasked implementation, specific instructions are used, while in our masked implementation, only a naive multiplication is applied. The problem is that most known optimized techniques require arithmetic operations, thus, a masked form would require a mask conversion. Given the complexity of such conversions, this approach may end up to be equivalent to our original naive technique. In the end, future work is still necessary to innovate and find new optimizations on this instruction.

Similarly, the cyclic shift is performed here directly, while the reference implementation stores the polynomials in duplicate (contiguously) and just has to change its "window" to perform the shift. We could not see any way to keep this advantage in a masked form. This also explains why there is such a difference in performance between the reference implementation and this implementation when the order equals 0.

General performances for masked BIKE (fully-dense)   The performances and scaling for the scheme are detailed in Table 5 and Fig. 8.

Remark 4. RNG off refers to returning 0 instead of drawing a random integer. This allows to measure the cost of the number of calls to the RNG, relative to the performance of the implementation.

We can see that the performance of masked BIKE as a function of the order is slightly above quadratic. This unoptimized implementation is still encouraging as it leaves the door open for many possible scaling improvements.
In fact, there are still a lot of possible optimizations, in particular on the cyclic shift and on the naive polynomial multiplication. Once optimized, the scaling will probably be improved, especially since there is no boolean arithmetic conversion within the masked scheme.

Table 5: Scaling benchmarks on BIKE, i7-4710MQ 2.5Ghz gcc 12.2.0 -03, NIST Level 1, median results on 100 executions, in million of cycles

|	Order 0	1	2	3	4	5
SecKeyGen (RNG off)	55	162	351	477	640	853
SecKeyGen (RNG on)	55	188	409	635	980	1 330
Scaling SecKeyGen (RNG off)	**1**	**×3**	**×6.4**	**×8.7**	**×11.7**	**×15.5**
Scaling SecKeyGen (RNG on)	**1**	**×3.4**	**×7.4**	**×11.5**	**×17.9**	**×24.2**
Encaps (RNG off)	5	24	53	84	120	170
Encaps (RNG on)	5	29	71	122	190	278
Scaling Encaps (RNG off)	**1**	**×4.8**	**×10.6**	**×16.8**	**×24**	**×34**
Scaling Encaps (RNG on)	**1**	**×5.8**	**×14.2**	**×24.4**	**×38**	**×55.6**
Decaps (RNG off)	63	262	559	842	1 220	1 652
Decaps (RNG on)	63	329	723	1 211	1 873	2 693
Scaling Decaps (RNG off)	**1**	**×4.1**	**×8.9**	**×13.4**	**×19.4**	**×26.2**
Scaling Decaps (RNG on)	**1**	**×5.2**	**×11.5**	**×19.2**	**×29.7**	**×42.7**
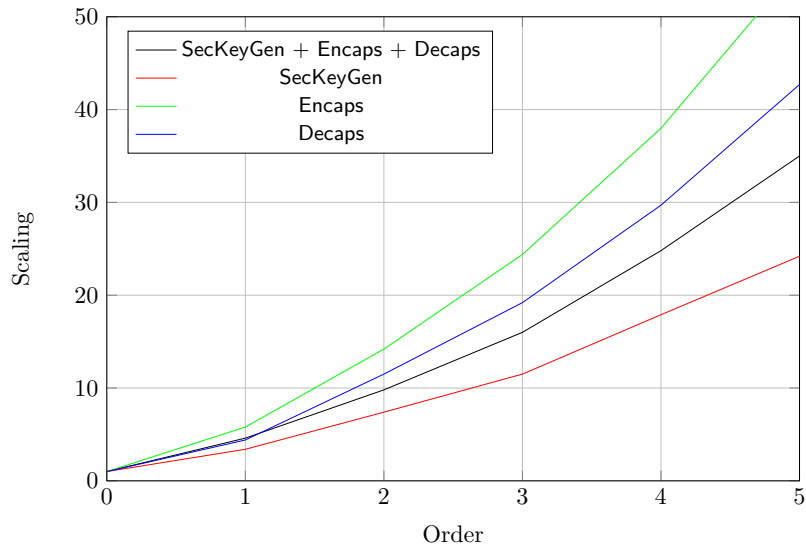
Figure 8: The scaling of masked BIKE (with RNG on)

6 Future Work

TVLA The next step would be to use TVLA verification techniques on our code to check that there are no apparent leaks.

More optimizations It is possible to highly optimize the performance of our implementation by simply optimizing two important basic gadgets: the naive multiplication (in the last level of the Karatsuba recursion) and the cyclic shift. As outlined above, these gadgets are the bottleneck of our implementation. Thus, the impact on the performance to be very high. The relevance of avoiding mask conversions may also be questioned if such conversions help to gain orders of magnitude in the performance; even though we do not currently believe that conversions would significantly help here. In addition, we think that further optimization could impact the difference between the sparse version and the dense version.

High-order attacks Attacking unprotected implementations with side-channel measurements is often not the best choice to evaluate practical security. But, until now, no masked implementation of BIKE and other code-based schemes were available. This masked implementation is openly accessible and can serve as target for elaborate high-order side-channel attacks.

References

[ABB+22] Nicolas Aragon, Paulo Barreto, Slim Bettaieb, Loic Bidoux, Olivier Blazy, Jean-Christophe Deneuville, Phillipe Gaborit, Shay Gueron, Tim Guneysu, Carlos Aguilar Melchor, Rafael Misoczki, Edoardo Persichetti, Nicolas Sendrier, Jean-Pierre Tillich, Gilles Zémor, Valentin Vasseur, Santosh Ghosh, and Jan Richter-Brokmann. BIKE. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/round-4-submissions.

[ABC+22] Melissa Azouaoui, Olivier Bronchain, Gaëtan Cassiers, Clément Hoffmann, Yulia Kuzovkova, Joost Renes, Markus Schönauer, Tobias Schneider, François-Xavier Standaert, and Christine van Vredendaal. Leveling dilithium against leakage: Revisited sensitivity analysis and improved implementations. Cryptology ePrint Archive, Paper 2022/1406, 2022. https://eprint.iacr.org/2022/1406.

[Ale03] Michael Alekhnovich. More on average case vs approximation complexity. pages 298–307, 2003.

[BAA+19] Nina Bindel, Sedat Akleylek, Erdem Alkim, Paulo S. L. M. Barreto, Johannes Buchmann, Edward Eaton, Gus Gutoski, Juliane Kramer, Patrick Longa, Harun Polat, Jefferson E. Ricardini, and Gustavo Zanon. qTESLA. Technical report, National Institute of Standards and Technology, 2019. available at https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions.

[BBD+16] Gilles Barthe, Sonia Belaïd, François Dupressoir, Pierre-Alain Fouque, Benjamin Grégoire, Pierre-Yves Strub, and Rébecca Zucchini. Strong non-interference and type-directed higher-order masking. pages 116–129, 2016.

[BBE+18] Gilles Barthe, Sonia Belaïd, Thomas Espitau, Pierre-Alain Fouque, Benjamin Grégoire, Mélissa Rossi, and Mehdi Tibouchi. Masking the GLP lattice-based signature scheme at any order. pages 354–384, 2018.

[BGR+21] Joppe W. Bos, Marc Gourjon, Joost Renes, Tobias Schneider, and Christine van Vredendaal. Masking kyber: First- and higher-order implementations. 2021(4):173–214, 2021. https://tches.iacr.org/index.php/TCHES/article/view/9064.

[BMvT78] E. Berlekamp, R. McEliece, and H. van Tilborg. On the inherent intractability of certain coding problems (corresp.). *IEEE Transactions on Information Theory*, 24(3):384–386, 1978.

[BOG20] Mario Bischof, Tobias Oder, and Tim Güneysu. Efficient Microcontroller Implementation of BIKE. In Emil Simion and Rémi Géraud-Stewart, editors, *Innovative Security Solutions for Information Technology and Communications*, pages 34–49, Cham, 2020. Springer International Publishing.

[CARG23] Agathe Cheriere, Nicolas Aragon, Tania Richmond, and Benoît Gérard. Bike key-recovery: Combining power consumption analysis and information-set decoding, 2023.

[CCK21] Ming-Shing Chen, Tung Chou, and Markus Krausz. Optimizing BIKE for the intel haswell and ARM cortex-M4. 2021(3):97–124, 2021. https://tches.iacr.org/index.php/TCHES/article/view/8969.

[CGKT22] Ming-Shing Chen, Tim Güneysu, Markus Krausz, and Jan Philipp Thoma. Carry-less to BIKE faster. pages 833–852, 2022.

[CGTV15] Jean-Sébastien Coron, Johann Großschädl, Mehdi Tibouchi, and Praveen Kumar Vadnala. Conversion from arithmetic to Boolean masking with logarithmic complexity. pages 130–149, 2015.

[CGV14] Jean-Sébastien Coron, Johann Großschädl, and Praveen Kumar Vadnala. Secure conversion between Boolean and arithmetic masking of any order. pages 188–205, 2014.

[Cho16] Tung Chou. QcBits: Constant-time small-key code-based cryptography. pages 280–300, 2016.

[CJRR99] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. pages 398–412, 1999.

[Cor14] Jean-Sébastien Coron. Higher order masking of look-up tables. pages 441–458, 2014.

[Cor17] Jean-Sébastien Coron. High-order conversion from Boolean to arithmetic masking. pages 93–114, 2017.

[CPRR14] Jean-Sébastien Coron, Emmanuel Prouff, Matthieu Rivain, and Thomas Roche. Higher-order side channel security and mask refreshing. pages 410–424, 2014.

[DDF14] Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: From probing attacks to noisy leakage. pages 423–440, 2014.

[DG17] Nir Drucker and Shay Gueron. A toolbox for software optimization of QC-MDPC code-based cryptosystems. Cryptology ePrint Archive, Report 2017/1251, 2017. https://eprint.iacr.org/2017/1251.

[DGK20] Nir Drucker, Shay Gueron, and Dusan Kostic. QC-MDPC decoders with several shades of gray. pages 35–50, 2020.

[DHP+21] Jan-Pieter D'Anvers, Daniel Heinz, Peter Pessl, Michiel van Beirendonck, and Ingrid Verbauwhede. Higher-order masked ciphertext comparison for lattice-based cryptography. Cryptology ePrint Archive, Report 2021/1422, 2021. https://eprint.iacr.org/2021/1422.

[DKR+20] Jan-Pieter D'Anvers, Angshuman Karmakar, Sujoy Sinha Roy, Frederik Vercauteren, Jose Maria Bermudo Mera, Michiel Van Beirendonck, and Andrea Basso. SABER. Technical report, National Institute of Standards and Technology, 2020. available at https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions.

[GAB19] Antonio Guimarães, Diego F. Aranha, and Edson Borin. Optimized implementation of QC-MDPC code-based cryptography. *Concurrency and Computation: Practice and Experience*, 31(18):e5089, 2019.

[Gal62] R. Gallager. Low-density parity-check codes. *IRE Transactions on Information Theory*, 8(1):21–28, 1962.

[GHJ+21]   Qian Guo, Clemens Hlauschek, Thomas Johansson, Norman Lahr, Alexander Nilsson, and Robin Leander Schröder. Dont́ reject this: Key-recovery timing attacks due to rejection-sampling in hqc and bike. Cryptology ePrint Archive, Paper 2021/1485, 2021. https://eprint.iacr.org/2021/1485.

[GJS16]   Qian Guo, Thomas Johansson, and Paul Stankovski. A key recovery attack on MDPC with CCA security using decoding errors. pages 789–815, 2016.

[GR19]   François Gérard and Mélissa Rossi. An efficient and provable masked implementation of qtesla. Cryptology ePrint Archive, Report 2019/606, 2019. https://eprint.iacr.org/2019/606.

[HHK17]   Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. pages 341–371, 2017.

[ISW03]   Yuval Ishai, Amit Sahai, and David Wagner. Private circuits: Securing hardware against probing attacks. pages 463–481, 2003.

[KDV+22]   Suparna Kundu, Jan-Pieter D'Anvers, Michiel Van Beirendonck, Angshuman Karmakar, and Ingrid Verbauwhede. Higher-order masked Saber. Cryptology ePrint Archive, Report 2022/389, 2022. https://eprint.iacr.org/2022/389.

[KLRBG22a] Markus Krausz, Georg Land, Jan Richter-Brockmann, and Tim Güneysu. Efficiently masking polynomial inversion at arbitrary order. Cryptology ePrint Archive, Report 2022/707, 2022. https://eprint.iacr.org/2022/707.

[KLRBG22b] Markus Krausz, Georg Land, Jan Richter-Brockmann, and Tim Güneysu. A holistic approach towards side-channel secure fixed-weight polynomial sampling. Cryptology ePrint Archive, Paper 2022/1740, 2022. https://eprint.iacr.org/2022/1740.

[LDK+22]   Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022.

[Lem19]   Daniel Lemire. Fast random integer generation in an interval. *ACM Transactions on Modeling and Computer Simulation*, 29(1):1–12, jan 2019.

[MOG15]   Ingo Von Maurich, Tobias Oder, and Tim Güneysu. Implementing QC-MDPC McEliece Encryption. *ACM Trans. Embed. Comput. Syst.*, 14(3), April 2015.

[MTSB13]   Rafael Misoczki, Jean-Pierre Tillich, Nicolas Sendrier, and Paulo S. L. M. Barreto. Mdpc-mceliece: New mceliece variants from moderate density parity-check codes. In *2013 IEEE International Symposium on Information Theory*, pages 2069–2073, 2013.

[Nie86]   H. Niederreiter. Knapsack-type cryptosystems and algebraic coding theory. *Problems of Control and Information Theory*, 15(2):159–166, 1986.

[RHHM17]   Melissa Rossi, Mike Hamburg, Michael Hutter, and Mark E. Marson. A side-channel assisted cryptanalytic attack against QcBits. pages 3–23, 2017.

[RP10]   Matthieu Rivain and Emmanuel Prouff. Provably secure higher-order masking of AES. pages 413–427, 2010.

[SAB+22]   Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, Damien Stehlé, and Jintai Ding. CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022.

[Sen21]   Nicolas Sendrier. Secure sampling of constant-weight words ? application to bike. Cryptology ePrint Archive, Report 2021/1631, 2021. https://eprint.iacr.org/2021/1631.

[SKC⁺19] Bo-Yeon Sim, Jihoon Kwon, Kyu Young Choi, Jihoon Cho, Aesun Park, and Dong-Guk Han. Novel side-channel attacks on quasi-cyclic code-based cryptography. 2019(4):180–212, 2019. https://tches.iacr.org/index.php/TCHES/article/view/8349.

[vMG14] Ingo von Maurich and Tim Güneysu. Towards side-channel resistant implementations of QC-MDPC McEliece encryption on constrained devices. pages 266–282, 2014.

[vMHG16] Ingo von Maurich, Lukas Heberle, and Tim Güneysu. IND-CCA secure hybrid encryption from QC-MDPC niederreiter. pages 1–17, 2016.

A Extra small gadgets and proofs

This part is about extra small gadgets and their proofs.

A.1 Adders and carries

Algorithm 23 Half Adder

Require: $[\![x]\!] \in \mathbb{F}_2$, $[\![y]\!] \in \mathbb{F}_2$
Ensure: $[\![z]\!] = [\![x]\!] + [\![y]\!] \in \mathbb{F}_2$, $[\![c]\!] \in \mathbb{F}_2$ the carry
1: $[\![z]\!] \leftarrow [\![x]\!] \oplus [\![y]\!]$
2: $[\![c]\!] \leftarrow \mathsf{sec}_\&([\![x]\!], [\![y]\!])$
3: **return** $[\![z]\!], [\![c]\!]$

Algorithm 24 Adder

Require: $[\![x]\!] \in \mathbb{F}_2$, $[\![y]\!] \in \mathbb{F}_2$, $[\![c_0]\!] \in \mathbb{F}_2$
Ensure: $[\![z]\!] = [\![x]\!] + [\![y]\!] + [\![c_0]\!] \in \mathbb{F}_2$, $[\![c]\!] \in \mathbb{F}_2$ the carry
1: $([\![t]\!], [\![s]\!]) \leftarrow \mathsf{SecHalf\_Adder}([\![x]\!], [\![y]\!])$
2: $([\![z]\!], [\![u]\!]) \leftarrow \mathsf{SecHalf\_Adder}([\![t]\!], [\![c_0]\!])$
3: $[\![c]\!] \leftarrow [\![s]\!] \oplus [\![u]\!]$
4: **return** $[\![z]\!], [\![c]\!]$

Since the $\oplus$ enjoys the d-NI property and $\mathsf{sec}_\&$ takes the same variables as input but is d-SNI, their combination leads to an d-NI algorithm. Thus, we introduce the following stating the probing security of the half adder algorithm.

Theorem 16. *The half adder algorithm is* $d - \mathsf{NI}$*.*

Since the adder uses only two calls to $\mathsf{SecHalf\_Adder}$ (itself $d - \mathsf{NI}$), handling different variables, we can infer the following.

Theorem 17. *The adder algorithm is* $d - \mathsf{NI}$*.*

A.2 Equality

The gadget $\mathsf{sec}_=$ is a d-NI gadget that outputs a masked Boolean value corresponding to the equality. The idea is to use Boolean algebra to check if the XOR between the two inputs is 0. For that, we perform a $\mathsf{sec}_\&$ between the negation of each obtained bit. Such a procedure has been outlined in the literature e.g. in [DHP⁺21].

We decided to optimize it in such a way as to dichotomize the operations about the word binary, and thus achieve better performance.

Algorithm 25 Masked equality (sec$_=$)

Require: $[\![\mathbf{x}]\!] \in \mathbb{F}_2^n, [\![\mathbf{y}]\!] \in \mathbb{F}_2^n, n$ a power of 2
Ensure: $[\![z]\!] \in \mathbb{F}_2$ equals 0 if $\mathbf{x} = \mathbf{y}$ and 1 if not
 1: $[\![\mathbf{z}]\!] \leftarrow [\![\mathbf{x}]\!] \oplus [\![\mathbf{y}]\!]$
 2: **for** $i \leftarrow \frac{n}{2}$ to 1 step $-\frac{i}{2}$ **do**
 3:      $[\![\mathbf{a}]\!] \leftarrow \mathsf{left}([\![\mathbf{z}]\!])$ ▷ Cut in length
 4:      $[\![\mathbf{b}]\!] \leftarrow \mathsf{right}([\![\mathbf{z}]\!])$ ▷ Cut in length
 5:      $[\![\mathbf{a}]\!]_0 \leftarrow \neg [\![\mathbf{a}]\!]_0$ ▷ Coefficient-wise not
 6:      $[\![\mathbf{b}]\!]_0 \leftarrow \neg [\![\mathbf{b}]\!]_0$ ▷ Coefficient-wise not
 7:      $[\![\mathbf{z}]\!] \leftarrow \mathsf{sec}_\&([\![\mathbf{a}]\!], [\![\mathbf{b}]\!])$ ▷ Coefficient-wise $\mathsf{sec}_\&$
 8:      $[\![\mathbf{z}]\!]_0 \leftarrow \neg [\![\mathbf{z}]\!]_0$ ▷ Coefficient-wise not
 9: **end for**
10: **return** $[\![\mathbf{z}_0]\!]$

Given that the only operation manipulating the data is $\mathsf{sec}_\&$, and that it is a d-SNI function, we can deduce that the algorithm is d-NI.

In fact, as negation only manipulates the first share, it is not able to leak anything (given that values are updated at each loop turn).

Theorem 18. *The equality algorithm is d-NI.*

A.3 Multiplications

Algorithm 26 Partly masked multiplication (SecMult$_\text{partlymasked}$)

Require: $[\![x]\!]$, y, $b = \lfloor \log_2(y) \rfloor$
Ensure: $[\![z]\!] = [\![x]\!] \cdot y$
 1: $[\![z]\!] \leftarrow \mathsf{zero\_masking}()$
 2: $[\![t]\!] \leftarrow [\![x]\!]$
 3: **for** $i \leftarrow 0$ to b **do**
 4:      **if** $y[i] = 1$ **then**   ▷ $y[i] = (y \gg i) \,\&\, 1$
 5:          $[\![z]\!] \leftarrow \mathsf{sec}_+([\![z]\!], [\![t]\!])$
 6:          $[\![t]\!] \leftarrow \mathsf{refresh}([\![t]\!])$
 7:      **end if**
 8:      $[\![t]\!] \leftarrow [\![t]\!] \ll 1$
 9: **end for**
10: **return** $[\![z]\!]$

Theorem 19. *The partly masked multiplication algorithm is $d - $ NI.*

Proof. There are two possible blocks in the for loop : if $y[i] = 0$, the only operation is the shift, which enjoys the d-NI property. If $y[i] = 1$, we use the sec_+ gadget which is also d-NI. Since we reuse the t in the shift, we need to refresh it before.

So the two blocks are d-NI, and their sequential combination leads to a d-NI algorithm   □

A.4 Modular random number

To generate keys and errors, we need to be able to draw random numbers modulo n.

For this, we are using a method formalized by Lemire [Lem19], which allows us to draw an integer between 0 and $n - 1$ with the same distribution as a modulo without performing any division other than with a power of 2. We will only need the gadgets already introduced (masked multiplication see Algorithm 26) and the shift, which is a linear operation.

Remark 5. It is assumed that the bits can be drawn safely, since the p bits can be drawn on each of the shares of the shared value. In the context of an implementation, the choice of algorithm for effectively drawing these bits is up to the developer.

Algorithm 27 Modular random number ($\mathsf{sec_{rand}}$)

Require: $n \in \mathbb{N}^*$, $p \in \mathbb{N}^*$, $2^p \geq n$

Ensure: $[\![r]\!] \stackrel{\$}{\leftarrow} \mathbb{Z}_n$

1: $[\![r]\!] \stackrel{\$}{\leftarrow} \mathbb{F}_{2^p}$ ▷ Draw p bits on each share
2: $[\![r]\!] \leftarrow \mathsf{SecMult}_{\mathrm{partlymasked}}([\![r]\!], n)$
3: $[\![r]\!] \leftarrow [\![r]\!] \gg p$ ▷ Shift on each share
4: **return** $[\![r]\!]$

Theorem 20. *The modular random number Algorithm 27 is d-NI.*

Proof. Since p and n are public values, we do not need to mask them.

Since it operates on each share individually, the shift operation is d-NI.

$\mathsf{SecMult}_{\mathrm{partlymasked}}$ is d-NI, by the previous proof. Finally, the random draw is also d-NI since it operates on each share.

The algorithm is d-NI.   □

Algorithm 28 SecPolymul: Naive Polynomial multiplication (parameterized by B, the size of its inputs)

Require: $[\![\mathbf{x}]\!] \in \mathbb{F}_2^B$, $[\![\mathbf{y}]\!] \in \mathbb{F}_2^B$

Ensure: $[\![\mathbf{z}]\!] = [\![\mathbf{x}]\!] \cdot [\![\mathbf{y}]\!] \in \mathbb{F}_2^{2B}$

1: **for** $i \leftarrow 0$ to $B - 1$ **do**
2:     **for** $j \leftarrow 0$ to $B - 1$ **do**
3:         $[\![u]\!] \leftarrow \mathsf{sec_\&}([\![\mathbf{x}_i]\!], [\![\mathbf{y}_j]\!])$
4:         $[\![\mathbf{z}_{(i+j)}]\!] \leftarrow [\![\mathbf{z}_{(i+j)}]\!] \oplus [\![u]\!]$
5:     **end for**
6: **end for**
7: **return** $[\![\mathbf{z}]\!]$

Since we only use a SNI gadget and we update the **z** vector on the other hand, the algorithm is $d - \mathsf{NI}$.

Theorem 21. *The polynomial multiplication SecPolymul parametered with B algorithm is $d - \mathsf{NI}$.*

A.5 Cyclic shift

This is a masked version of the barrel shifter algorithm.

We define SecCyclic_Shift the function that allows to shift a masked polynomial with a public value. As it is only a linear operation, it is safe and not a concern.

Algorithm 29 Secure masked cyclic shift ($\mathsf{sec_\gg}$)

Require: $[\![\mathbf{x}]\!] \in \mathbb{F}_2^n$, $[\![s]\!] \in \mathbb{N}$

Ensure: $[\![\mathbf{y}]\!] = [\![\mathbf{x}]\!] \gg [\![s]\!] \in \mathbb{F}_2^n$

1: $[\![\mathbf{y}]\!] \leftarrow [\![\mathbf{x}]\!]$
2: **for** $i \leftarrow 0$ to $|n|$ **do**   ▷ $|n| = \lfloor \log_2(n) \rfloor$
3:     $[\![v]\!] \leftarrow [\![s]\!][i]$ ▷ $[\![s]\!][i] = ([\![s]\!] \gg i) \mathbin{\&} 1$
4:     $[\![\mathbf{t}]\!] \leftarrow \mathsf{SecCyclic\_Shift}([\![\mathbf{y}]\!], 2^i)$
5:     **for** $j \leftarrow 0$ to $n - 1$ **do**
6:         $[\![s1]\!] \leftarrow \mathsf{sec_\&}([\![\mathbf{t}_j]\!], [\![v]\!])$
7:         $[\![s2]\!] \leftarrow \mathsf{sec_\&}([\![\mathbf{y}_j]\!], \neg[\![v]\!])$
8:         $[\![\mathbf{y}_j]\!] \leftarrow [\![s1]\!] \oplus [\![s2]\!]$
9:     **end for**
10: **end for**
11: **return** $[\![\mathbf{y}]\!]$

Theorem 22. *The secure cyclic shift algorithm is $d - \mathsf{NI}$.*

Proof. In the most imbricated for loop, we used two $\mathsf{sec}_\&$, which are d-SNI. The $\oplus$ being d-NI, the block is d-NI. Since the i loop is composed by d-NI gadgets, and the $\mathbf{y}$ vector is updated in the for j loop, all of this is d-NI. So their sequential combination leads to a d-NI algorithm. $\qquad\square$

A.6 Masked conditional branch

Algorithm 30 Choose value ($\mathsf{sec}_{\mathsf{if}}$)

Require: $[\![a]\!] \in \mathbb{F}_n$, $[\![b]\!] \in \mathbb{F}_n$, $[\![t]\!] \in \mathbb{F}_2$
Ensure: $[\![a]\!]$ if $[\![t]\!] = 1$, $[\![b]\!]$ otherwise
1: $[\![c]\!] \leftarrow \mathsf{sec}_\&^{\mathbf{bitwise}}([\![a]\!], [\![t]\!])$ $\triangleright$ Bitwise $\mathsf{sec}_\&$ between all bits of a and the single bit of t
2: $[\![t]\!]_0 \leftarrow \neg [\![t]\!]_0$
3: $[\![d]\!] \leftarrow \mathsf{sec}_\&^{\mathbf{bitwise}}([\![b]\!], [\![t]\!])$ $\triangleright$ Bitwise $\mathsf{sec}_\&$ between all bits of b and the single bit of t
4: **return** $[\![c]\!] \oplus [\![d]\!]$ $\triangleright$ Coefficient-wise XOR

Since $\mathsf{sec}_\&^{\mathbf{bitwise}}$ is just a succession of $\mathsf{sec}_\&$ which is d-SNI, $\mathsf{sec}_\&^{\mathbf{bitwise}}$ is also d-SNI property. Since the last $\oplus$ is d-NI, we deduce the theorem below.

Theorem 23. *The choose value algorithm is $d - \mathsf{NI}$.*

A.7 Masked maximum computation

Algorithm 31 Max ($\mathsf{sec}_{\mathsf{max}}$)

Require: $[\![a]\!] \in \mathbb{Z}_n$, $[\![b]\!] \in \mathbb{Z}_n$
Ensure: $[\![c]\!] = \mathsf{sec}_{\mathsf{max}}([\![a]\!], [\![b]\!]) \in \mathbb{Z}_n$
1: $[\![t]\!] \leftarrow \mathsf{sec}_+([\![a]\!], [\![-b]\!])$
2: **return** $\mathsf{sec}_{\mathsf{if}}(\mathsf{refresh}([\![b]\!]), \mathsf{refresh}([\![a]\!]), \mathsf{sign\_bit}([\![t]\!]))$

Since the variables a and b are used within the sec_+ gadget, which is d-NI, we need to refresh them (d-SNI gadget) before reusing them in the call to the $\mathsf{sec}_{\mathsf{if}}$ function. This yields the following theorem.

Theorem 24. *The max algorithm is $d - \mathsf{NI}$.*

A.8 Filling a matrix in masked form

Algorithm 32 Fill matrix ($\mathsf{sec}_{\mathsf{fill}}$)

Require: $[\![v]\!] \in \mathbb{Z}_n$
Ensure: $[\![\mathbf{X}]\!] \in \mathbb{F}_2^{k \times (|n|+1)}$ a matrix filled with the binary representation of $[\![v]\!]$
1: **for** $i \leftarrow 0$ to $k - 1$ **do**
2:     **for** $j \leftarrow 0$ to $|n|$ **do**
3:         $[\![\mathbf{X}_{i,j}]\!] \leftarrow [\![v]\!][j]$
4:     **end for**
5:     $[\![v]\!] \leftarrow \mathsf{refresh}([\![v]\!])$
6: **end for**
7: **return** $[\![\mathbf{X}]\!]$

Since we just initialize $[\![\mathbf{X}]\!]$ with $[\![v]\!]$ binary, we just refresh $[\![v]\!]$ to avoid to get same mask in two different lines.
We then get the following theorem.

Theorem 25. *The fill algorithm is $d - \mathsf{NI}$.*