Round5:
KEM and PKE based on GLWR
Draft Friday 25th January, 2019

Hayo Baan[1], Sauvik Bhattacharya[1], Scott Fluhrer[5], Oscar Garcia-Morchon[1], Thijs Laarhoven[3], Ludo Tolhuizen[1], Ronald Rietman[1], Markku-Juhani O. Saarinen[2], and Zhenfei Zhang[4]

[1]Philips (Netherlands)
[2]PQShield (UK)
[3]TU/e (Netherlands)
[4]Algorand (US)
[5]Cisco (US)

Abstract

Standardization bodies such as NIST and ETSI are currently seeking quantum resistant alternatives to vulnerable RSA and elliptic curve-based public-key algorithms. In this context, we present Round5, a lattice-based cryptosystem providing a key encapsulation mechanism and a public-key encryption scheme. Round5 is based on the General Learning with Rounding problem, unifying non-ring and ring lattice rounding problems into one. Usage of rounding combined with a tight analysis leads to significantly reduced bandwidth and randomness requirements. Round5's reliance on prime-order cyclotomic rings offers a large design space allowing fine-grained parameter optimization. The use of sparse-ternary secret keys improves performance and significantly reduces decryption failure rates at minimal additional cost. The use of error-correcting codes, in combination with ring multiplications in $\mathbb{Z}[x]/(x^{n+1}-1)$ that ensures non-correlated errors, further improves the latter. Round5 parameters have been carefully optimized for bandwidth, while the design facilitates efficient implementation.

As a result, Round5 has leading performance characteristics among all NIST post-quantum candidates, and at the same time attains conservative security levels that fully fit NIST's security categories. Round5's schemes share common building blocks, simplifying (security and operational) analysis and code review. Finally, Round5 proposes various approaches of refreshing the system public parameter A, which efficiently prevent precomputation and back-door attacks.

Disclaimer: This is a draft version, not all sections are included.

Contents

1   Algorithm Specifications

1.1   Design Rationale

This submission proposes Round5 that consists of algorithms for the key encapsulation mechanism r5_cpa_kem and the public-key encryption scheme r5_cca_pke. The proposed algorithms fall under the category of lattice-based cryptography, Round5 is a merger of the submissions Round2[13, 14] and HILA5[85, 86]. Like with Round2, the algorithms rely on the General Learning with Rounding (GLWR) problem. In some of its configurations, Round5 uses an error-correcting code based on the one from HILA5 to decrease decryption failure probability, and thus, achieve smaller keys and faster performance.

1.1.1   A Unified Design

A key feature of Round5 is that it has been designed to instantiate the Learning with Rounding (LWR) problem and the Ring LWR (RLWR) problem in a seamless and unified way. This is done by defining the General LWR problem, on which Round5 is based, that can instantiate LWR or RLWR depending on the input parameters. The reasons behind this choice are as follows:

Round5 is adaptive and can be applied to multiple environments. On the one hand, LWR-based algorithms are required by environments in which performance is less of an issue, but security is the priority. In those cases, it is often preferred to not have an additional ring structure (as in ideal lattices [70, 54]). On the other hand, RLWR-based algorithms achieve the best performance in terms of bandwidth and computation so they are better suited for constrained environments with stricter bandwidth requirements, e.g., due to the complexity of message fragmentation or small MTUs.

Round5 reduces code analysis and maintenance since the unified scheme definitions of r5_cpa_kem and r5_cca_pke instantiate different underlying problems, LWR and RLWR, with the same code.

The unified design enables a migration strategy from ring-based schemes to non-ring schemes from day one of deployment. This makes sure that if vulnerabilities in ring-based problems were found in future, then an alternative secure solution would already be available and could be deployed directly.

An additional advantage of a unified design is that it can be "parametrized" to address the needs of different types of applications. This brings an enormous flexibility. For instance, it is possible to derive specially small configurations applicable to IoT scenarios or it is also possible to derive non-ring configurations with balanced or unbalanced public-key and ciphertext sizes.

1.1.2   Parameter choices for optimized performance

Parameters in GLWR and Round5 are chosen to allow for optimized performance.

- The usage of GLWR, i.e., LWR and RLWR, rather than their LWE counterparts, leads to lower bandwidth requirements in Round5 since fewer bits need to be transmitted per coefficient.

- Rounding avoids sampling of noise, and thus requires the generation of less random data.

- Round5 relies on a definition of GLWR with sparse ternary secrets. This simplifies implementation and reduces the probability of decryption/decapsulation errors.

- The ring configuration of Round5 relies on the RLWR problem over the cyclotomic ring $\mathbb{Z}_q[x]/\Phi_{n+1}(x)$ with $n + 1$ prime. This problem is well studied: there exist reductions [18, 23] from the RLWE problem [70] to RLWR, and the former is well-studied for the ring in question. Operations over this ring can be mapped to an NTRU [54] ring $\mathbb{Z}_q[x]/(x^{n+1} - 1)$ to improve performance.

- In some of its configurations, Round5 uses an f-bit error correcting block code XEf to decrease the failure rate, see Section 1.4.1. The code is built using the same strategy as codes used by TRUNC8 [86] (2-bit correction) and HILA5 [87] (5-bit correction). The main advantage of XEf codes is that they avoid table look-ups and conditions altogether and are therefore resistant to timing attacks. The usage of XEf requires performing the ciphertext computations in the NTRU ring and sparse ternary secrets that are balanced, i.e., contain equally many ones as minus ones. This leads to independent bit failures so that error correction can be applied (Section 1.4.1).

- Round5 allows for a single, *unified* implementation for a wide range of security levels, relying on either LWR or RLWR. The moduli q and p are powers of two. This simplifies the implementation of the rounding function. Similarly, the modular computations can be realized by ignoring the most significant bits.

- Preventing pre-computation attacks requires refreshing the master public parameter $\boldsymbol{A}$. However, this can be computationally costly, in particular in the case of LWR. Round5 provides several alternatives for efficiently refreshing $\boldsymbol{A}$ that are applicable to different use cases.

1.1.3   The Choice of the Ring

In the literature, a common choice of the ring to instantiate an RLWE or RLWR problem is $\mathbb{Z}_q[x]/\Phi_{2n}(x)$ where n is a power of two, so that the $2n$-th cyclotomic polynomial $\Phi_{2n}(x) = x^n + 1$. Examples of RLWE/RLWR key exchange schemes based on the above ring are [28] and [7]. However, requiring that n be a power of two narrows the choice of n: it has to be at least 512 for the proper security level so that the underlying lattice problem is hard to solve in practice. While having

$n = 512$ sometimes does not deliver the target security level, the $n = 1024$ choice would be considered an overkill. A sweet spot is $n \approx 700$, which was a common choice made by many proposals, including Kyber [26], NTRUEncrypt [54], NTRU-KEM [58] and more.

The following observations can be made:

- Kyber [26] uses three RLWE instances, each of dimension 256, to achieve $n = 768$ in total. This still limits the choice of n as it has to be a multiple of 256.

- NTRUEncrypt uses the reduction polynomial $x^n - 1$ which is slightly different from a cyclotomic ring, and as suggested by [84], although the NTRU problem remains hard for this ring, the decisional RLWE problem over this ring seems to be easy.

This leads us to use as reduction polynomial the $n + 1$-th cyclotomic polynomial $\Phi_{n+1}(x) = x^n + \cdots + x + 1$ with $n + 1$ a prime as in NTRU-KEM [58]. With the resulting ring, there is a wide range of n to choose from, for various security levels. In addition, as shown in [78], decisional RLWE over this ring remains hard for any modulus; this gives us confidence in the underlying design and security of Round5. Note that although operating over the same ring as the NTRU-KEM scheme, Round5 achieves better performance because its key generation algorithm is significantly faster than the NTRU key generation algorithm.

In addition, since decisional RLWE is hard over a prime cyclotomic ring with any modulus [78], we can use any modulus that optimizes our performances. Common choices of the modulus are

- A number theoretical transform (NTT) friendly prime number, such as 12289 in [7];

- A composite number that fits in a data type for modern computers, such as $2^{32} - 1$ in [28];

- A power of two that makes modulo operations and integer multiplications efficient, such as 2^{11} in NTRUEncrypt [54].

In our proposal, we consider a prime cyclotomic polynomial ring with a modulus q that is a power of two such that the Φ_{n+1} is irreducible modulo two. As Φ_{n+1} then is irreducible modulo q, the ring $\mathbb{Z}_q[x]/\Phi_{n+1}(x)$ does not have any proper subrings. This choice allows for a smooth implementation of the ring and non-ring cases in the unified scheme. All coefficients of our polynomials and matrices are integers modulo q less than 2^{16}. That is, each coefficient fits in a `uint16_t` type. For intermediate values during computations, overflows can be ignored, as overflown bits "*mod-ed out*" once the element is lifted back to $\mathbb{Z}_q$. In particular, when multiplying two `uint16_t` elements, only the lower 16 bits of the product need to be computed; the higher 16 bits have no effect on the final result.

We remark that the use of a composite modulus in [28], as well as the result from [78] suggest that the particular modulus does not have much effect on the hardness of the problem; the size of the modulus is more important from this point of view. Consequently, any modulus of similar size should deliver a similar security, and therefore we choose one that is most efficient depending on the use case.

1.2   Preliminaries

For each positive integer a, we denote the set $\{0, 1, \ldots, a-1\}$ by $\mathbb{Z}_a$.

For a set A, we denote by $a \xleftarrow{\$} A$ that a is drawn uniformly from A. If χ is a probability distribution, then $a \leftarrow \chi$ means that a is drawn at random according to the probability distribution χ.

Logarithms are in base 2, unless specified otherwise.

Modular reductions. For a positive integer α and $x \in \mathbb{Q}$, we define $\{x\}_\alpha$ as the unique element x' in the interval $(-\alpha/2, \alpha/2]$ satisfying $x' \equiv x \pmod{\alpha}$. Moreover, we define $\langle x \rangle_\alpha$ as the unique element x' in the interval $[0, \alpha)$ for which $x \equiv x' \pmod{\alpha}$.

Rounding. For $x \in \mathbb{Q}$, we denote by $\lfloor x \rfloor$ and $\lfloor x \rceil$ rounding downwards to the next smaller integer and rounding to the closest integer (with rounding up in case of a tie) respectively.

For positive integers a, b and $h \in \mathbb{Q}$, the rounding function $\mathrm{R}_{a \to b, h}$ is defined as

$$\mathrm{R}_{a \to b, h}(x) = \langle \lfloor \frac{b}{a}(x+h) \rfloor \rangle_b \tag{1}$$

If $h = \frac{a}{2b}$, then $\mathrm{R}_{a \to b, h}(x) = \langle \lfloor \frac{b}{a} x \rceil \rangle_b$. This special case of rounding will be used in the underlying problem, so the following notation will come in handy: for positive integers a, b, the function $\mathrm{R}_{a \to b}$ is defined as

$$\mathrm{R}_{a \to b} = \mathrm{R}_{a \to b, a/2b} \tag{2}$$

Ring choice. Let $n+1$ be prime. The $(n+1)$-th cyclotomic polynomial $\Phi_{n+1}(x)$ then equals $x^n + x^{n-1} + \cdots + x + 1$. We denote the polynomial ring $\mathbb{Z}[x]/\Phi_{n+1}(x)$ by $\mathcal{R}_n$. When n equals 1, then $\mathcal{R}_n = \mathbb{Z}$. For each positive integer a, we write $\mathcal{R}_{n,a}$ for the set of polynomials of degree less than n with all coefficients in $\mathbb{Z}_a$. We call a polynomial in $\mathcal{R}_n$ *ternary* if all its coefficients are 0, 1 or -1. Throughout this document, regular font letters denote elements from $\mathcal{R}_n$, and bold lower case letters represent vectors with coefficients in $\mathcal{R}_n$. All vectors are column vectors. Bold upper case letters are matrices. The transpose of a vector $\boldsymbol{v}$ or a matrix $\boldsymbol{A}$ is denoted by $\boldsymbol{v}^T$ or $\boldsymbol{A}^T$. A vector or matrix of polynomials, in which all polynomial entries are equal to 1 is denoted by $\boldsymbol{j}$ or $\boldsymbol{J}$, respectively.

Distributions. For each $v \in \mathcal{R}_n$, the **Hamming weight** of v is defined as its number of non-zero coefficients. The Hamming weight of a vector in $\mathcal{R}_n^k$ equals the sum of the Hamming weighs of its components. We denote with $\mathcal{H}_{n,k}(h)$ the set of all vectors $\boldsymbol{v} \in \mathcal{R}_n^k$ of ternary polynomials of Hamming weight h, where $h \leq nk$. By considering the coefficients of a polynomial in $\mathcal{R}_n$ as a vector of length n, a polynomial in $\mathcal{H}_{n,k}(h)$ corresponds to a ternary vector of length nk with non-zeros in h positions, so that $\mathcal{H}_{n,k}(h)$ has $\binom{nk}{h}2^h$ elements. When $k = 1$, we omit it from the notation, and $\mathcal{H}_n(h)$ denotes the set of all ternary polynomials in $\mathcal{R}_n$ of Hamming weight h, corresponding to the set of all vectors $\boldsymbol{v} \in \{-1, 0, 1\}^n$ with Hamming weight h.

Secret keys consist of matrices that contain (column) vectors in $\mathcal{H}_{n,k}(h)$. Functions f_R and f_S are used to generate secrets from a seed in the encryption (Algorithm 2) and decryption (Algorithm 3), respectively.

1.3 Underlying Problem

The problem underlying the security of Round5 is the **General LWR Problem** formally defined as follows:

Definition 1.3.0.1 (General LWR (GLWR)). *Let d, n, p, q be positive integers such that $q \geq p \geq 2$, and $n \in \{1, d\}$. Let $\mathcal{R}_{n,q}$ be a polynomial ring, and let D_s be a probability distribution on $\mathcal{R}_n^{d/n}$.*

The search version of the GLWR problem $\mathtt{sGLWR}_{d,n,m,q,p}(D_s)$ is as follows: given m samples of the form $R_{q \to p}(\boldsymbol{a}_i^T \boldsymbol{s})$ with $\boldsymbol{a}_i \in \mathcal{R}_{n,q}^{d/n}$ and a fixed $\boldsymbol{s} \leftarrow D_s$, recover $\boldsymbol{s}$.

The decision version of the GLWR problem $\mathtt{dGLWR}_{d,n,m,q,p}(D_s)$ is to distinguish between the uniform distribution on $\mathcal{R}_{n,q}^{d/n} \times \mathcal{R}_{n,p}$ and the distribution $(\boldsymbol{a}_i, b_i = R_{q \to p}(\boldsymbol{a}_i^T \boldsymbol{s}))$ with $\boldsymbol{a}_i \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n}$ and a fixed $\boldsymbol{s} \leftarrow D_s$.

When $n = 1$, the GLWR problem is equivalent to the LWR problem [18] with dimension d, large modulus q and rounding modulus p. Setting the distribution $D_s = \mathcal{U}(\mathcal{H}_{1,d}(h))$ further specializes the GLWR problem to the LWR problem with sparse-ternary secrets $\mathtt{LWR}_{\text{spt}}$. In [35] it is claimed that the hardness of $\mathtt{LWR}_{\text{spt}}$ can be obtained from that of LWE with similar secret distributions since the reduction from LWE to LWR is independent of the secret's distribution [23]. We extend this claim and make it explicit by proving that for appropriate parameters there exists a polynomial-time reduction from the (decision) Learning with Errors (LWE) problem with secrets chosen uniformly from $\mathbb{Z}_q^d$ and errors chosen from a Gaussian distribution D_α, to the decision version of $\mathtt{LWR}_{\text{spt}}$. See Section A.8 and Theorem A.8.0.1 for more details.

When $n = d > 1$ is such that $n + 1$ is prime, and $\mathcal{R}_{n,q} = \mathbb{Z}_q[x]/(\Phi_{n+1}(x))$ for the $n + 1$-th cyclotomic polynomial $\Phi_{n+1}(x) = 1 + x + \ldots + x^n$, the GLWR problem is equivalent to the Ring LWR (RLWR) problem defined on $\Phi_{d+1}(x)$, dimension d, large modulus q and rounding modulus p. Setting $D_s = \mathcal{U}(\mathcal{H}_{d,1}(h))$ further specializes it to the RLWR problem with sparse-ternary secrets $\mathtt{RLWR}_{\text{spt}}$. For brevity, when $D_s = \mathcal{U}(\mathcal{H}_{n,d/n}(h))$, we denote $\mathtt{GLWR}_{d,n,m,q,p}(\mathcal{U}(\mathcal{H}_{n,d/n}(h)))$

as $\texttt{GLWR}_{\text{spt}}$. When the secret distribution D_s is the uniform one over $\mathcal{R}_{n,q}^{d/n}$, it shall be omitted in the above problem notation.

1.4 Round5

Our proposal is called Round5. It includes an IND-CPA secure KEM called r5_cpa_kem and a IND-CCA secure PKE called r5_cca_pke. A public-key encryption scheme called r5_cpa_pke is used as a building block for r5_cpa_kem. A key-encapsulation mechanism called r5_cca_kem, is used as a building block for r5_cca_pke. All algorithms in these schemes use random choices, and scheme-specific mappings that are described in the following sections with each scheme.

In Section 1.4.1, the error-correcting codes applied in some configurations of Round5 are described. These codes are built using the same strategy as codes used by TRUNC8 [86] (2-bit correction) and HILA5 [87] (5-bit correction). r5_cpa_kem is described in Section 1.4.3, preceded by the description of its building block r5_cpa_pke in Section 1.4.2. r5_cca_pke is descirbed in Section 1.4.5, preceded by its building block r5_cca_kem in Section 1.4.4. Figure 1 provides an overview of our proposal. It also shows the different configurations of the proposed schemes based on the underlying GLWR problem.
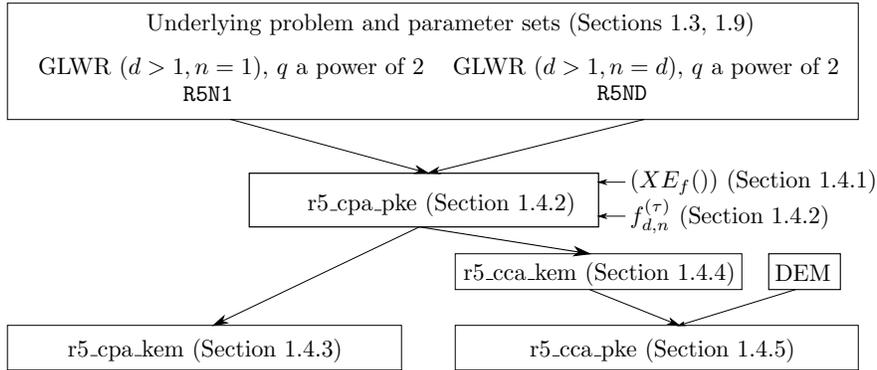
Figure 1: Submission overview

1.4.1 Internal building block: error correction code

In [46], it is analyzed how error-correcting codes can be used the enhance the error resilience of protocols like NewHope, Frodo and Kyber, and it is shown that the usage of error correcting codes can significantly increase the estimated bit-security and decrease the communication overhead. In some of its configurations, Round5 uses an f-bit error correcting block code XEf to decrease the failure rate. The code is built using the same strategy as codes used by TRUNC8 [86] (2-bit correction) and HILA5 [87] (5-bit correction).

XEf. The XEf code is described by $2f$ "registers" r_i of size $|r_i| = l_i$ with $i = 0, \ldots, 2f - 1$. We view the κ-bits payload block m as a binary polynomial $m_{\kappa-1}x^{\kappa-1} + \cdots + m_1 x + m_0$ of length κ. Registers are defined via cyclic reduction

$$r_i = m \bmod x^{l_i} - 1, \tag{3}$$

or equivalently by

$$r_i[j] = \sum_{k \equiv j \bmod l_i} m_k \tag{4}$$

where $r_i[j]$ is bit j of register r_i. A transmitted message consists of the payload m concatenated with register set r (a total of $\mu = \kappa + xe$ bits, where $xe = \sum l_i$).

Upon receiving a message $(m' \mid r')$, one computes the register set r'' corresponding to m' and compares it to the received register set r' – that may also have errors. Errors are in coefficients m'_k where there are parity disagreements $r'_i[k \bmod l_i] \neq r''_i[k \bmod l_i]$ for multitude of registers r_i. We use a majority rule and flip bit m'_k if

$$\sum_{i=0}^{2f-1} ((r'_i[\langle k \rangle_{l_i}] - r''_i[\langle k \rangle_{l_i}]) \bmod 2) \geq f + 1 \tag{5}$$

where the sum is taken as the number of disagreeing register parity bits at k.

XEf codes can include a special register – that is marked with symbol $(*)$ such that $\ell^{(*)}$ divides κ, and the parity bits in register $r^{(*)}$ are computed as the sum of $\kappa/\ell^{(*)}$ consecutive bits. That is, for $0 \leq j \leq l^{(*)} - 1$, bit $r^{(*)}[j]$ is defined as

$$r^{(*)}[j] = \sum_{i=0}^{\frac{\kappa}{\ell^{(*)}}-1} m_{j \cdot \frac{\kappa}{\ell^{(*)}} + i}.$$

In HILA5[87]'s XE5 code, r_0 is such a special register. When using this special register, decoding is done as explained above, but with the term corresponding to the special register $r^{(*)}$ replaced by $\left(r'^{(*)}[\lfloor \frac{k\ell^{(*)}}{\kappa} \rfloor] - r''^{(*)}[\lfloor \frac{k\ell^{(*)}}{\kappa} \rfloor] \right) \bmod 2$.

As shown in HILA5[87], if all length pairs satisfy $\text{lcm}(l_i, l_j) \geq \kappa$ when $i \neq j$, then this code always corrects at least f errors. If the XEf code includes a special register $r^{(*)}$, say $r^{(*)} = r_0$, then it is required that $\frac{\kappa}{\ell_0} \leq l_i$ for all $j \geq 1$ and $\text{lcm}(l_i, l_j) \geq \kappa$ for $i, j \geq 1$ and $i \neq j$.

The main advantage of XEf codes is that they avoid table look-ups and conditions altogether and are therefore resistant to timing attacks.

Requirements for using XEf in Round5. A basic requirement for using XEf error correction code is that the errors it aims to correct are independent. As explained in Section 1.4.2, Round5 can use two reduction polynomials $\xi(x) = \Phi_{n+1}(x) = \frac{x^{n+1}-1}{x-1}$ or $\xi(x) = N_{n+1}(x) = x^{n+1} - 1$ in the computation of the ciphertext and in the decryption/decapsulation. As explained in more detail in Section 1.8, using $\xi(x) = \Phi_{n+1}(x)$ leads to correlated errors. The basic reason is that the difference polynomial $d(x)$ governing the error behavior is obtained

as $d(x) = \langle \sum_{i=0}^{n} f_i x^i \rangle_{\Phi_{n+1}(x)} = \sum_{i=0}^{n} f_i x^i - f_n \Phi_{n+1}(x) = \sum_{i=0}^{n-1} (f_i - f_n) x^i$. As a result, if f_n is large, then many coefficients of $d(x)$ are likely to be large, leading to errors in the corresponding bit positions. As a consequence of this correlation between errors, the XEf code cannot be directly employed with the reduction polynomial $\xi(x) = \Phi_{n+1}(x)$ as used in Round2.

As detailed in Section 1.8, Round5 aims to obtain independent errors by imposing two requirements:

- Ciphertext computation and decryption use the reduction polynomial $\xi(x) = N_{n+1}(x)$.

- The secret ternary polynomials $s(x)$ and $r(x)$ are balanced, that is, both have as many coefficients equal to 1 as to -1.

The latter requirement implies that $(x-1)$ is a factor of both $s(x)$ and $r(x)$. As a consequence, for any polynomial $a(x)$, it holds that $\langle s(x)a(x)\rangle_{\Phi_{n+1}(x)} r(x) \equiv s(x)\rangle a(x)r(x)\rangle_{\Phi_{n+1}(x)} \pmod{N_{n+1}(x)}$. This equivalence relation is essential for operational correctness, cf. Section 1.8.

These two requirements are aligned with features of the original Round2 submission [13] since (i) Round2 already internally performed all operations in $\xi(x) = N_{n+1}(x)$ and (ii) Round2's implementation used balanced secrets.

1.4.2   Internal building block: r5_cpa_pke

This section describes r5_cpa_pke, the CPA-secure public key encryption that is a building block in both r5_cpa_kem and r5_cca_pke. It consists of algorithms 1 (key-generation), 2 (encryption) and 3 (decryption), and various cryptosystem parameters, *viz* positive integers n, d, h, p, q, t, b, $\overline{n}$, $\overline{m}$, μ, f, τ, and a security parameter κ. The system parameters are listed in Table 1. In the proposed configurations, $n \in \{1, d\}$, and b, q, p, t are powers of 2, such that $b|t|p|q$. It is required that

$$\mu \leq \overline{n} \cdot \overline{m} \cdot n \text{ and } \mu * b\_bits \geq \kappa \text{ where } b = 2^{b\_bits}.$$

The function $f_{d,n}^{(\tau)}$ generates a public matrix $\boldsymbol{A}$ from a seed σ. Round5 has three options for $f_{d,n}^{(\tau)}$:

- $f_{d,n}^{(0)}$ generates $\boldsymbol{A}$ by means of a deterministic random bit generator (DRBG) from a seed σ specific for the protocol exchange. This is a standard approach followed by many other protocols such as Frodo, Kyber, or NewHope.

- $f_{d,n}^{(1)}$ generates $\boldsymbol{A}$ by permuting a public and system-wide matrix $\boldsymbol{A_{master}}$. The permutation is derived from a seed σ. With this approach, it is possible to efficiently obtain a fresh $\boldsymbol{A}$ that depends on a seed σ – so that backdoor and pre-computation attacks are avoided – without having to compute lots of pseudorandom data, so that higher efficiency is achieved.

Table 1: List of symbols and their meaning

n	Indicates if ring ($n > 1$) or ring-configuration ($n = 1$) is used
d	Dimension of underlying lattice
h	Number of non-zero components per column of secret matrices
b, p, q, t	Rounding moduli, all powers of two, satisfying $b < t < p < q$
b_bits	Number of extracted bits per symbol in ciphertext component; $b = 2^{b\_bits}$.
$\overline{n}$	Number of columns of the secret matrix of the initiator
$\overline{m}$	Number of columns of the secret matrix of the responder
κ	Security parameter; number of information bits in error-correcting code
xe	Number of parity bits of error correcting code
μ	Number of symbols in ciphertext component: $\mu = \lceil \frac{\kappa + xe}{b_{bits}} \rceil$.
Sample_μ	Function for picking up μ polynomial coefficients from a matrix
f	Number of bit errors correctable by error-correcting code
τ	Index for method of determining public matrix from a seed in $\{0,1\}^\kappa$
$f_{d,n}^{(\tau)}$	Method for determining public matrix from a seed in $\{0,1\}^\kappa$
f_S	Function for determining secret matrix for initiator from a seed in $\{0,1\}^\kappa$
f_R	Function for determining secret matrix for responder from a seed in $\{0,1\}^\kappa$
z	$z = \max(p, \frac{tq}{p})$. Relevant in reduction proofs and definition of rounding constants
h_1	$h_1 = \frac{q}{2p}$ is a rounding constant
h_2	$h_2 = \frac{q}{2z}$ is a rounding constant
h_3	$h_3 = \frac{p}{2t} + \frac{p}{2b} - \frac{q}{2z}$ is a constant for reducing bias in decryption
h_4	$h_4 = \frac{q}{2p} - \frac{q}{2z}$ is a constant for reducing bias in decryption
$\Phi_{n+1}(x)$	$\Phi_{n+1}(x) = \frac{x^{n+1} - 1}{x - 1}$. Reduction polynomial for computation of public keys; $n+1$ a prime.
$\xi(x)$	Reduction polynomial for computation of ciphertext. $\xi(x) \in \{\Phi_{n+1}(x), x^{n+1} - 1\}$

- $f_{d,n}^{(2)}$ generates $\boldsymbol{A}$ by first generating a vector $\boldsymbol{a}$ from a seed σ and then permuting it. The permutation is also random and derived from seed σ as well. With this approach, it is possible to efficiently obtain a fresh $\boldsymbol{A}$ that depends on a seed σ – so that backdoor and pre-computation attacks are avoided – without having to compute lots of pseudorandom data and without having to store a potentially large matrix $\boldsymbol{A_{master}}$.

In the rest of this section, options like the choice of the DRBG underlying $f_{d,n}^{(\tau)}$ are hidden. The specific implementation of $f_{d,n}^\tau$ and performance trade-offs discussed in Section 1.11.4 and Section 1.9

The function $\text{Sample}_\mu : \boldsymbol{C} \in \mathcal{R}_{n,p}^{\overline{n} \times \overline{m}} \to \mathbb{Z}_p^\mu$ outputs μ polynomial coefficients of the $\overline{n} \cdot \overline{m} \cdot n$ polynomial coefficients present in $\boldsymbol{C}$. These μ coefficients are used in the actual encryption algorithm. This function is fully specified in

Section 1.11.4.

The integer f denotes the error-correction capability of a code $XE_{\kappa,f} \subset \mathbb{Z}_b^{\mu}$. We have an encoding function $xef\_compute_{\kappa,f} : \{0,1\}^{\kappa} \to XE_{\kappa,f}$ and a decoding function $xef\_decode_{\kappa,f} : \mathbb{Z}_b^{\mu} \to \{0,1\}^{\kappa}$ such that for each $m \in \{0,1\}^{\kappa}$ and each error $e = (e_0, \ldots, e_{\mu-1})$ with at most f symbols e_i different from zero,

$$xef\_decode_{\kappa,f}(xef\_compute_{\kappa,f}(m) + e) = m. \qquad (6)$$

The functions $xef\_compute()$ and $xef\_decode()$, based on Xef codes, are detailed from an implementation perspective in Section 1.11.4.

Algorithm r5_cpa_pke_encrypt employs a deterministic function f_R for generating a secret matrix $\boldsymbol{R} \in \left(\mathcal{H}_{n,d/n}(h)\right)^{1 \times \overline{m}}$ from an input ρ. Defining ρ as an explicit input to r5_cpa_pke_encrypt allows us to reuse this *same* algorithm as a building block for both IND-CPA and IND-CCA secure cryptographic constructions.

Furthermore, r5_cpa_pke uses four rounding constants, namely

$$h_1 = \frac{q}{2p}, \; h_2 = \frac{q}{2z}, h_3 = \frac{p}{2t} + \frac{p}{2b} - \frac{q}{2z} \text{ and } h_4 = \frac{q}{2p} - \frac{q}{2z}, \text{ with } z = \max(p, \frac{tq}{p})$$
$$\qquad (7)$$

The constant h_1 leads to rounding to the closest integer. The choice of h_2 ensures that Round5's ciphertext $(\boldsymbol{U}, \boldsymbol{v})$ is provably pseudorandom under the GLWR assumption. Details are provided in the proof of IND-CPA security for r5_cpa_pke and r5_cpa_kem (Section A.4). The choices of h_3 and h_4 are made to avoid bias in decryption, see Section 1.8.

Note that the rounding constants are added explicitly here for completeness, but with the specific parameter choices in the different Round5 configurations 1.4.6 some of them can be simplified: $h_1 = h_2 = q/2p$ leading to standard rounding as in Round2 [13, 14] implemented by means of flooring. Furthermore, $h_4 = 0$. Thus, the only difference with respect to Round2 is h_3, which is present to avoid bias in decryption.

Algorithm 1: r5_cpa_pke_keygen()

---

    parameters: Integers $p, q, n, h, d, \overline{n}, \kappa, \tau$
    input   : -
    output   : $pk \in \{0,1\}^{\kappa} \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}, sk \in \{0,1\}^{\kappa}$

1 $\sigma \xleftarrow{\$} \{0,1\}^{\kappa}$

2 $\boldsymbol{A} = \boldsymbol{f}_{d,n}^{(\tau)}(\sigma)$

3 $sk \xleftarrow{\$} \{0,1\}^{\kappa}$

4 $\boldsymbol{S} = f_S(sk)$

5 $\boldsymbol{B} = \mathrm{R}_{q \to p, h_1}(\langle \boldsymbol{AS} \rangle_{\Phi_{n+1}})$

6 $pk = (\sigma, \boldsymbol{B})$

7   **return** (pk, sk)

Algorithm 2: r5_cpa_pke_encrypt(pk, m, ρ)

---

    parameters: Integers $p, t, q, n, d, \overline{m}, \overline{n}, \mu, b, \kappa, f, \tau; \xi \in \{\Phi_{n+1}(x), x^{n+1} - 1\}$
    input   : $pk = (\sigma, \boldsymbol{B}) \in \{0,1\}^{\kappa} \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}, m, \rho \in \{0,1\}^{\kappa}$
    output   : $ct = (\boldsymbol{U}, \boldsymbol{v}) \in \mathcal{R}_{n,p}^{d/n \times \overline{m}} \times \mathbb{Z}_t^{\mu}$

1 $\boldsymbol{A} = \boldsymbol{f}_{d,n}^{(\tau)}(\sigma)$

2 $\boldsymbol{R} = f_R(\rho)$

3 $\boldsymbol{U} = \mathrm{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$

4 $\boldsymbol{v} = \langle \mathrm{R}_{p \to t, h_2}(\mathrm{Sample}_{\mu}(\langle \boldsymbol{B}^T \boldsymbol{R} \rangle_{\xi})) + \frac{t}{b} xef\_compute_{\kappa,f}(m) \rangle_t$

5 $ct = (\boldsymbol{U}, \boldsymbol{v})$

6   **return** ct

Algorithm 3: r5_cpa_pke_decrypt(sk, ct)

---

    parameters: Integers $p, t, q, n, d, \overline{m}, \overline{n}, \mu, b, \kappa, f; \xi \in \{\Phi_{n+1}(x), x^{n+1} - 1\}$
    input   : $sk \in \{0,1\}^{\kappa}, ct = (\boldsymbol{U}, \boldsymbol{v}) \in \mathcal{R}_{n,p}^{d/n \times \overline{m}} \times \mathbb{Z}_t^{\mu}$
    output   : $\hat{m} \in \{0,1\}^{\kappa}$

1 $\boldsymbol{v}_p = \frac{p}{t}\boldsymbol{v}$

2 $\boldsymbol{S} = f_s(sk)$

3 $\boldsymbol{y} = \mathrm{R}_{p \to b, h_3}(\boldsymbol{v}_p - \mathrm{Sample}_{\mu}((\boldsymbol{S}^T(\boldsymbol{U} + h_4 \boldsymbol{J}))_{\xi}))$

4 $\hat{m} = xef\_decode_{\kappa,f}(\boldsymbol{y})$

5   **return** $\hat{m}$

1.4.3   Submission proposal: r5_cpa_kem

This section describes r5_cpa_kem, an IND-CPA-secure key encapsulation method. It builds on r5_cpa_pke (Section 1.4.2). In addition to the parameters and functions from r5_cpa_pke, it uses a hash function $H : \{0,1\}^* \to \{0,1\}^{\kappa}$. In order to improve readability, in Algorithms 5 and 6 the conversion of the ciphertext ct into a binary string before it is fed to H is not made explicit.

Algorithm 4: r5_cpa_kem_keygen()

---

    parameters: Integers $p, q, n, h, d, \overline{n}, \kappa, \tau$
    input   : -
    output   : $pk \in \{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}, sk \in \{0,1\}^\kappa$

1 $(pk, sk) = $ r5_cpa_pke_keygen()
2 **return** (pk, sk)

Algorithm 5: r5_cpa_kem_encapsulate(pk)

---

    parameters: Integers $p, t, q, n, d, \overline{m}, \overline{n}, \mu, b, \kappa, f, \tau; \xi \in \{\Phi_{n+1}(x), x^{n+1} - 1\}$
    input   : $pk \in \{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}$
    output   : $(ct, K) \in (\mathcal{R}_{n,p}^{d/n \times \overline{m}} \times \mathbb{Z}_t^\mu) \times \{0,1\}^\kappa$

1 $m \xleftarrow{\$} \{0,1\}^\kappa$
2 $\rho \xleftarrow{\$} \{0,1\}^\kappa$
3 $ct = $ r5_cpa_pke_encrypt(pk, m, ρ)
4 $K = H(m, ct)$
5 **return** (ct, K)

Algorithm 6: r5_cpa_kem_decapsulate(sk, ct)

---

    parameters: Integers $p, t, q, n, d, \overline{m}, \overline{n}, \mu, b, \kappa, f; \xi \in \{\Phi_{n+1}(x), x^{n+1} - 1\}$
    input   : $sk \in \{0,1\}^\kappa, ct \in \mathcal{R}_{n,p}^{d/n \times \overline{m}} \times \mathbb{Z}_t^\mu$
    output   : $K \in \{0,1\}^\kappa$

1 $m = $ r5_cpa_pke.decrypt(sk, ct)
2 $K = H(m, ct)$
3 **return** K

1.4.4   Internal building block: r5_cca_kem

This section describes r5_cca_kem that is a building block for r5_cca_pke. It consists of the algorithms 7, 8, 9, and several system parameters and functions in addition to those from r5_cpa_pke and r5_cpa_kem. In addition to the hash function H from r5_cpa_kem, it uses another hash function $G : \{0,1\}^* \rightarrow \{0,1\}^\kappa \times \{0,1\}^\kappa \times \{0,1\}^\kappa$. To improve readability, conversion of data types to bitstrings before being fed to G and H is not made explicit.

    r5_cca_kem is actively secure as it is obtained by application of the Fujisaki-Okamoto transform [56] to r5_cpa_pke, similarly as in [27, Sec. 4]. On decapsulation failure, i.e. if the condition in line 4 of Algorithm 9 is not satisfied, a pseudorandom key is returned, causing later protocol steps to fail implicitly. Explicit failure notification would complicate analysis, especially in the quantum random oracle (QROM) case.

Algorithm 7: r5_cca_kem_keygen()

---

parameters: Integers $p, q, n, h, d, \overline{n}, \kappa, \tau$
input   : -
output   : $pk \in \{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}, sk \in \{0,1\}^\kappa \times \{0,1\}^\kappa \times (\{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}})$
1 $(pk, sk_{CPA-PKE}) = $ r5_cpa_pke_keygen()
2 $y \xleftarrow{\$} \{0,1\}^\kappa$
3 $sk = (sk_{CPA-PKE}, y, pk)$
4 **return** (pk, sk)

Algorithm 8: r5_cca_kem_encapsulate(pk)

---

parameters: Integers $p, t, q, n, d, \overline{m}, \overline{n}, \mu, b, \kappa, f, \tau; \xi \in \{\Phi_{n+1}(x), x^{n+1} - 1\}$
input   : $pk \in \{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}$
output   : $ct = (\boldsymbol{U}, \boldsymbol{v}, g) \in \mathcal{R}_{n,p}^{d/n \times \overline{m}} \times \mathbb{Z}_t^\mu \times \{0,1\}^\kappa, K \in \{0,1\}^\kappa$
1 $m \xleftarrow{\$} \{0,1\}^\kappa$
2 $(L, g, \rho) = G(m, pk)$
3 $(\boldsymbol{U}, \boldsymbol{v}) = $ r5_cpa_pke_encrypt(pk, m, ρ)
4 $ct = (\boldsymbol{U}, \boldsymbol{v}, g)$
5 $K = H(L, ct)$
6 **return** (ct, K)

Algorithm 9: r5_cca_kem_decapsulate(sk, ct)

---

parameters: Integers $p, t, q, n, d, \overline{m}, \overline{n}, \mu, b, \kappa, f, \tau; \xi \in \{\Phi_{n+1}(x), x^{n+1} - 1\}$
input   : $sk = (sk_{CPA-PKE}, y, pk) \in \{0,1\}^\kappa \times \{0,1\}^\kappa \times (\{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}), \;\; ct = (\boldsymbol{U}, \boldsymbol{v}, g) \in \mathcal{R}_{n,p}^{d/n \times \overline{m}} \times \mathbb{Z}_t^\mu \times \{0,1\}^\kappa$
output   : $K \in \{0,1\}^\kappa$
1 $m' = $ r5_cpa_pke_decrypt($sk_{CPA-PKE}, (\boldsymbol{U}, \boldsymbol{v})$)
2 $(L', g', \rho') = G(m', pk)$
3 $(\boldsymbol{U'}, \boldsymbol{v'}) = $ r5_cpa_pke_encrypt(pk, m', ρ')
4 **if** $(\boldsymbol{U'}, \boldsymbol{v'}, g') = (\boldsymbol{U}, \boldsymbol{v}, g)$ **then**
5   **return** $K = H(L', \boldsymbol{U}, \boldsymbol{v}, g)$
6 **else**
7   **return** $K = H(y, \boldsymbol{U}, \boldsymbol{v}, g)$
8 **end if**

1.4.5 Submission proposal: r5_cca_pke

The IND-CCA [82] public key encryption scheme r5_cca_pke consists of algorithms 10, 11 and 12. It combines r5_cca_kem with a data encapsulation mechanism (DEM), in the canonical way proposed by Cramer and Shoup [38]. r5_cca_kem is used to encapsulate a key K that is then used by the DEM to encrypt an arbitrary-length plaintext, optionally adding integrity protection. In decryption, r5_cca_kem is used to decapsulate K, which is then used by the DEM to decrypt and authenticate the plaintext.

Algorithm 10: r5_cca_pke_keygen()

---

parameters: Integers $p, q, n, h, d, \overline{n}, \kappa, \tau$

input : -

output : $pk \in \{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}, sk \in \{0,1\}^\kappa \times \{0,1\}^\kappa \times (\{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}})$

1 $(pk, sk) = $ r5_cca_kem_keygen()

2 return (pk, sk)

Algorithm 11: r5_cca_pke_encrypt(pk, M)

---

parameters: Integers $p, t, q, n, d, \overline{m}, \overline{n}, \mu, b, \kappa, f, \tau; \xi \in \{\Phi_{n+1}(x), x^{n+1} - 1\}$

input : $pk \in \{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}, mlen \in \mathbb{Z}, M \in \mathbb{Z}_{256}^{mlen}$

output : $ct = (c1, clen, c2) \in (\mathcal{R}_{n,p}^{d/n \times \overline{m}} \times Z_t^\mu \times \{0,1\}^\kappa) \times \mathbb{Z} \times \mathbb{Z}_{256}^{clen}$

1 $(c1, K) = $ r5_cca_kem_encapsulate(pk)

2 $(clen, c2) = $ DEM(K, M)

3 $ct = (c1, clen, c2)$

4 return ct

Algorithm 12: r5_cca_pke_decrypt(sk, ct)

---

parameters: Integers $p, t, q, n, d, \overline{m}, \overline{n}, \mu, b, \kappa, f, \tau; \xi \in \{\Phi_{n+1}(x), x^{n+1} - 1\}$

input : $sk \in \{0,1\}^\kappa \times \{0,1\}^\kappa \times (\{0,1\}^\kappa \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}), ct = (c1, clen, c2) \in (\mathcal{R}_{n,p}^{d/n \times \overline{m}} \times Z_t^\mu \times \{0,1\}^\kappa) \times \mathbb{Z} \times \mathbb{Z}_{256}^{clen}$

output : $M \in \mathbb{Z}_{256}^{mlen}$

1 $K = $ r5_cca_kem_decapsulate($sk, c1$)

2 $(mlen, M) = $ DEM$^{-1}(K, c2)$

3 return $(mlen, M)$

1.4.6 Proposed configurations

As illustrated by Table 1, Round5 has a large design space so that it can be easily configured to fit the needs of different applications. In particular, Round5 optimizes over the above design space to define the following configurations:

- Six ring configurations (r5_cpa_kem and r5_cca_pke for NIST security levels 1,3 and 5) without error correction, see Tables 2 and 11. These parameter choices can be considered to be conservative, as they are only based on the Round2 design that has received public review since its submission.

- Six ring configurations (r5_cpa_kem and r5_cca_pke for NIST security levels 1,3 and 5) with XEf error correction code, see Tables 3 and 12. Using XEf requires that $\xi(x) = x^{n+1} - 1$ and that the sparse ternary secrets are balanced (cf. Section 1.4.1. These parameter choices are based on the merge of HILA5 with Round2 and lead to the smallest public key and ciphertext sizes.

- Six non-ring configurations (r5_cpa_kem and r5_cca_pke for NIST security levels 1,3 and 5) without error correction, see Tables 4 and 13 These parameter choices rely on same design choices as the original Round2

submission. In particular, it uses $\overline{n} \approx \overline{m}$ to minimize total size of public-key plus ciphertext.

Round5 further details three additional specific use-case parameters with the only purpose of demonstrating its flexibility:

- A ring-based KEM configuration, addressing IoT. This configuration has lower bandwidth and computational requirements,yet still providing 90 bits of (quantum) security. XE2 forward error correction is used to improved failure rate, bandwidth requirements and security. See Table 5.

- A ring-based NIST level 1 configuration with Xef code in which the encapsulated key is 192-bit long instead of just 128-bit so that the difficulty of attacking the encapsulated key (by Grover) equals the difficulty of quantum lattice attack to Round5, see Table 5.

- A non-ring-based PKE parameter set with NIST level 3 with a ciphertext size of only 988 Bytes, with very fast encryption and decryption, by taking $\overline{m} = 1$, at the cost of a larger public key, see Table 14. This configuration targets applications in which the public-key can remain static for a long period, e.g., a fixed VPN between two end points. In such applications, the bandwidth footprint depends on the size of the ciphertext. Hence, this parameter set, despite enjoying the more conservative security assumptions for unstructured lattices, has a bandwidth requirement comparable to ring variants.

In contrast to the original Round2 and HILA5 submissions, Round5 does not include parameter sets suitable for NTT. The reason is that non-NTT parameters allow achieving better CPU-performance, as the advantages of using modular arithmetic with powers of two outweigh the advantages of using an NTT.
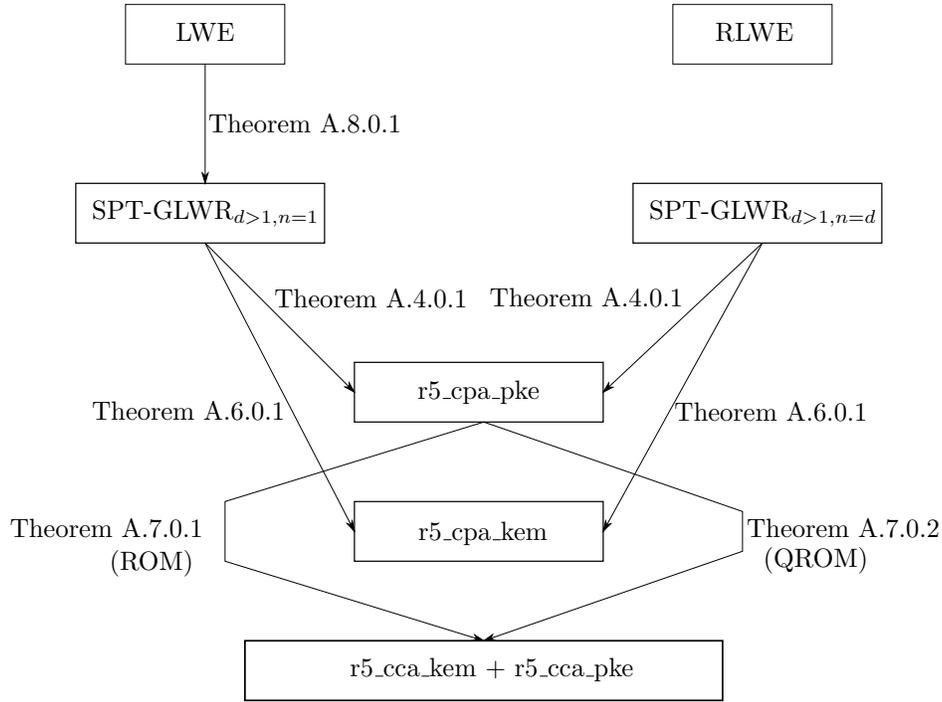
Figure 2: Summary of reductions involved in the security proofs for Round5 algorithms. "SPT" refers to a variant of the problem in question where the secret is sparse-ternary, instead of uniform in $\mathbb{Z}_q^d$.

1.5 Known Answer Test Values

The Known Answer Test Values for all algorithm variants and NIST levels can be found on the digital media at the location described in Section **??**.

Note that for generating intermediate output, the code needs to be compiled with -DROUND5_INTERMEDIATE (this option is enabled by default when making use of the provided Makefiles).

1.6 Expected Security Strength

This section summarizes results on the expected security of Round5 algorithms, its building blocks and underlying hard problem. An overview of these results is also given in Figure 2. For proofs, we refer to Appendix A.

1. **r5_cpa_pke** is an IND-CPA secure public-key encryption scheme if the decision General Learning with Rounding (GLWR) problem with sparse - ternary secrets is hard for the polynomial ring $\mathbb{Z}[x]/\Phi_{n+1}(x)$. Theorem A.4.0.1 gives a tight, classical reduction against classical or quantum adversaries in the standard model, a concise version of which is as follows:

Theorem 1.6.0.1. *For every adversary $\mathcal{A}$, there exist distinguishers $\mathcal{B}$, $\mathcal{C}$, $\mathcal{D}$, $\mathcal{E}$, $\mathcal{F}$ such that*

$$
\begin{aligned}
Adv_{CPA\text{-}PKE}^{IND\text{-}CPA}(\mathcal{A}) &\leq Adv_{\mathcal{U}(\mathcal{R}_{n,q}^{d/n \times d/n})}^{f_{n,d}^{\tau}}(\mathcal{B}) + Adv_{\chi_S^{\overline{n}}}^{f_S(s)}(\mathcal{C}) + \\
Adv_{\chi_S^{\overline{m}}}^{f_R(\rho)}(\mathcal{D}) &+ \overline{n} \cdot Adv_{d,n,d/n,q,p}^{dGLWR_{spt}}(\mathcal{E}) + Adv_{d,n,d/n+\overline{n},q,z}^{dGLWR_{spt}}(\mathcal{F})
\end{aligned}
\tag{8}
$$

where $z = \max(p, tq/p)$.

The full proof of IND-CPA security, given in Appendix A.4, follows a similar approach as [40] to equalize the noise ratios q/p and p/t in the two ciphertext components $\boldsymbol{U}$ and $\boldsymbol{v}$ that allows their combination into a single GLWR sample with noise ratio q/z. This proof does not apply if the reduction polynomial $\xi(x)$ in Round5 is $N_{n+1}(x)$, since then $\boldsymbol{U}$ and $\boldsymbol{v}$ cannot be combined into a single GLWR sample.

2. Appendix A.5 presents a proof of IND-CPA security for a Ring LWE based variant of r5_cpa_pke with reduction polynomial $\xi(x) = N_{n+1}(x)$, if the decision Ring LWE problem for $\mathbb{Z}[x]/\Phi_{n+1}(x)$ is hard; this results gives confidence for the RLWR case. Theorem A.5.0.1 gives a tight, classical reduction against classical and quantum adversaries in the standard model, a concise version of which is as follows:

Theorem 1.6.0.2. *For every adversary $\mathcal{A}$, there exist distinguishers $\mathcal{C}$ and $\mathcal{E}$ such that*

$$
Adv_{CPA\text{-}PKE, N_{n+1}(x)}^{IND\text{-}CPA}(\mathcal{A}) \leq Adv_{m=1}^{RLWE(\mathbb{Z}[x]/\phi(x))}(\mathcal{C}) + Adv_{m=2}^{RLWE(\mathbb{Z}[x]/\phi(x))}(\mathcal{E}).
\tag{9}
$$

where m denotes the number of available RLWE samples available.

This section also gives indications on sufficient conditions for the proof to work. One of them is replacing coordinate-wise rounding as done in Round5 by a more complicated form of rounding which ensures that the sum of the errors induced by rounding sum to zero. We chose not to use this form of rounding as it adds complexity and seems not to improve practical security. Moreover, the Sample$_\mu$ function in Round5 stops a well- known distinguishing attack against schemes based on rings with reduction polynomial $x^{n+1} - 1$.

3. **r5_cpa_kem** is an IND-CPA secure KEM if r5_cpa_pke is an IND-CPA secure public-key encryption scheme, and that H is a secure pseudo-random function. Theorem A.6.0.1 gives a tight, classical reduction against classical or quantum adversaries in the standard model; the construction and proof technique are standard.

4. **r5_cca_pke** is constructed from r5_cca_kem and a one-time data encapsulation mechanism in the canonical way proposed by Cramer and Shoup [39]. Therefore, if r5_cca_kem is IND-CCA secure, and the data encapsulation

mechanism is (one-time) secure against chosen ciphertext attacks and has a keylength fitting the security level of r5_cca_kem, then r5_cca_pke is an IND-CCA secure PKE. Section A.7 contains details.

5. **r5_cca_kem** is constructed using a KEM variant of the Fujisaki-Okamoto transform [53] from r5_cpa_pke. Therefore, assuming that r5_cpa_pke is an IND-CPA secure public-key encryption scheme, and G and H are modeled as random oracles, r5_cca_kem is an IND-CCA secure KEM. Theorem A.7.0.1 gives a tight, classical reduction against classical adversaries in the classical random oracle model. Theorem A.7.0.2 gives a non-tight, classical reduction against quantum adversaries in the quantum random oracle model.

6. Appendix A.1 gives an overview of the security reduction when replacing the GLWR public parameter $\boldsymbol{A}$ sampled from a truly uniform distribution with one generated in a pseudorandom fashion using the function $f_{d,n}^{(\tau)}$, for $\tau \in \{0, 1\}$. The reduction models AES(128 or 256) as an ideal cipher, and SHAKE(128 or 256) as a random oracle.

7. The decision Learning with Rounding (LWR) problem with sparse-ternary secrets is hard if the decision Learning with Errors (LWE) problem with uniform secrets and Gaussian errors is hard. Theorem A.8.0.1 gives a classical reduction against classical or quantum adversaries in the standard model, under the condition that the noise rate in the LWE problem decreases linearly in the security parameter n of the LWE and LWR problems. A concise version of this theorem is presented below:

Theorem 1.6.0.3. *Let p divide q, and $k \geq \frac{q}{p} \cdot m$. Let $\epsilon \in (0, \frac{1}{2})$, and $\alpha, \delta > 0$ such that*

$$\alpha \geq q^{-1}\sqrt{(2/\pi)\ln(2n(1+\epsilon^{-1}))}, \ \binom{n}{h}2^h \geq q^{k+1} \cdot \delta^{-2}, \ and \ m = O(\frac{\log n}{\alpha\sqrt{10h}})$$

Then there is a reduction from $\mathtt{dLWE}_{n,\frac{q}{p},q,D_{\alpha\sqrt{10h}}}\left(\mathcal{U}(\mathcal{H}_n\,(h))\right)$ *to* $\mathtt{dLWR}_{n,m,q,p}\left(\mathcal{U}(\mathcal{H}_n\,(h))\right)$.

1.7 Analysis with respect to known attacks

In this section, we analyze the concrete security of Round5. As the lattice-based attacks all rely on estimating the costs of solving certain lattice problems, we start with preliminaries on this topic in Sections 1.7.1 and 1.7.2. We then consider attacks using lattice basis reduction in Sections 1.7.3, 1.7.4 and 1.7.5, followed by specialized attacks that exploit sparse-ternary secrets used in Round5 in Sections 1.7.6 and 1.7.7. Finally, in Section 1.7.8 we consider precomputation and back-door attacks against the Round5 GLWR public parameter $\boldsymbol{A}$.

1.7.1 Preliminaries: SVP Complexities

On arbitrary lattices, we define the Shortest Vector Problem (SVP) as the problem of finding a vector in a lattice whose Euclidean norm is minimal among all non-zero vectors in the lattice. This problem is known to be NP-hard for randomized reductions [2, 62], and we do not expect any polynomial time algorithms for solving this problem in the worst-case to be found any time soon.

Various algorithms have been studied for solving SVP both in the worst-case and for average-case instances. The algorithm with the current best worst-case (asymptotic) complexity for solving SVP is the discrete Gaussian sampling approach [1], with an asymptotic cost in dimension n of $2^{n+o(n)}$ time and space. For large dimensions, this beats other exponential time and space algorithms based on computing the Voronoi cell [93, 72] or running a lattice sieve [3, 81], and enumeration algorithms requiring superexponential time in the lattice dimension [80, 61, 92, 47, 11]. These worst-case analyses however seem to suffer heavily from the (potential) existence of exotic, highly dense lattices such as the Leech lattice [36], and for average-case instances one can often solve SVP much faster, as also shown by various experiments on random lattices [41].

By making heuristic assumptions on the expected average-case behavior of SVP algorithms on random lattices, various works have demonstrated that some algorithms can solve SVP much faster than the above worst-case bounds suggest. For a long time, lattice enumeration [47, 11] was considered the fastest method in practice for solving SVP in high dimensions, and the crossover point with sieving-based approaches appeared to be far out of reach, i.e. well beyond cryptographically relevant parameters [73]. Recently, improvements to lattice sieving have significantly weakened this view, and the current fastest implementation based on sieving [4] has surpassed the best previous enumeration implementation both in terms of speed and in terms of the highest dimension reachable in practice. Although sieving is hindered by its exponential memory footprint, sieving also scales better in high dimensions compared to enumeration in terms of the time complexity ($2^{\Theta(n)}$ for sieving versus $2^{\Theta(n \log n)}$ for enumeration).

Conservative estimates. To estimate the actual computational complexity of solving SVP, we will use the best asymptotics for sieving [19], which state that solving SVP in dimension n takes time $2^{0.292n+o(n)}$. Clearly the hidden order term in the exponent has a major impact on the actual costs of these methods, and one might therefore be tempted to choose this term according to the actual, experimental value. However, recent improvements have shown that although the leading term $0.292n$ in the exponent is unlikely to be improved upon any time soon [10, 52], the hidden order term may well still be improved with heuristic tweaks such as those discussed in [42, 66, 4]. A conservative and practical cost estimate is therefore to estimate the costs of solving SVP in dimension n as $2^{0.292n}$, ignoring the (positive) hidden order term which may still be reduced over the next decades. Observe that this estimate is also well below the lowest estimates for enumeration or any other SVP method in high

dimensions.

Note that some work has also suggested that SVP is slightly easier to solve on structured, ideal lattices when using lattice sieving [60, 30, 20, 95]. However, these are all improvements reducing the time and/or space complexities by a small polynomial factor in the lattice dimension. Moreover, in our cost estimates we will actually be using an SVP algorithm as a subroutine within a blockwise lattice reduction algorithm, and even if the original lattice on which we run this lattice basis reduction algorithm has additional structure, these low-dimensional sublattices do not. Therefore, even for ideal lattices we do not expect these polynomial speedups to play a role.

Quantum speedups. For making Round5 post-quantum secure, we also take into account any potential quantum speedups to attacks an adversary might run against our scheme. Concretely, for the hardness of SVP on arbitrary lattices this may reduce the cost of lattice sieving to only $2^{0.265n+o(n)}$, assuming the attacker can efficiently run a Grover search [48] on a dynamically changing, exponential-sized database of lattice vectors stored in quantum RAM [68, 65]. Although it may not be possible to ever carry out such an attack efficiently, a conservative estimate for the quantum hardness of SVP in dimension n is therefore to assume this takes the attacker at least $2^{0.265n}$ time [65].

Note that recent work has also indicated that enumeration may benefit from a quantum speedup when applying a Grover-like technique to improve the backtracking search of lattice enumeration [12]. Disregarding significant polynomial overhead as well as simply the overhead of having to do operations quantumly, a highly optimistic estimate for the time costs of quantum enumeration in dimension n is $2^{\frac{1}{2}(0.187n\log n - 1.019n + 16.1)}$ [59]. Even with this model, the attack costs for all our parameter sets are higher than $2^{0.265n}$.

1.7.2 Lattice basis reduction: the core model

The above discussion focused on algorithms for solving exact SVP in high dimensions. The concrete security of lattice-based cryptographic schemes like Round5 however relies on the hardness of unique/approximate SVP, where a break constitutes finding a vector not much longer than the actual shortest non-zero vector in the lattice. This problem has also been studied extensively, and the the fastest method known to date for approximate SVP is the BKZ algorithm [91, 89, 32, 74, 4]. Given an arbitrary basis of a lattice, this algorithm (a generalization of the LLL algorithm [69]) performs HKZ reduction [63] on (projected) parts of the basis of a certain block size b. To perform HKZ reduction, the algorithm makes calls to an exact SVP oracle on b-dimensional lattices, and uses the result to improve the basis. Both the quality and time complexity can be tuned by b: running BKZ with larger b gives better bases, but also takes longer to finish.

Although the exact overall time complexity of BKZ remains somewhat mysterious [32, 51, 74, 4], it is clear that the dominant cost of BKZ for large block sizes is running the SVP oracle on b-dimensional lattices. The BKZ algorithm

further makes tours through the entire basis on overlapping blocks of basis vectors, but the number of tours only contributes a small multiplicative term, and SVP calls on overlapping blocks do not necessarily add up to independent SVP calls - for multiple overlapping blocks, it may well be possible to solve SVP in almost the same time as solving only one instance of SVP when using a sieving-based SVP solver [4].

This all gives rise to the core SVP model, where the complexity of BKZ with block size b is modeled by just *one* call to an SVP oracle in a b-dimensional lattice. This is clearly a lower bound for the actual cost of BKZ, and as further advances may be made to sieving-based BKZ algorithms reusing information between blocks, this is also the largest lower bound we can use without risking that predicted future advances in cryptanalysis will reduce the security level of our scheme in the near future.

1.7.3 Lattice-based attacks

We consider lattice-reduction based attacks, namely, the *primal* or decoding attack [16] and the *dual* or distinguishing attack [5], and how they can be adapted to exploit the shortness of secrets in our schemes. We begin by detailing how an attack on Round5 can be formulated as a lattice-reduction based attack on the LWR problem. We then analyze the concrete security of Round5 against the primal attack in order to estimate secure parameters, in Section 1.7.4. We do the same for the dual attack in Section 1.7.5.

The attacker can use the public keys $\boldsymbol{B} = \langle \lfloor \frac{p}{q} \langle \langle \boldsymbol{AS} \rangle_q \rangle_q \rceil \rangle_p$ of the public-key encryption scheme or the key-encapsulation scheme to obtain information on the secret key $\boldsymbol{S}$. We work out how this is done. For the non-ring case, $\boldsymbol{B} \in \mathbb{Z}_p^{d \times \overline{n}}$. Note also that since $q|p$ in this case, $\boldsymbol{B} = \langle \lfloor \frac{p}{q} \langle \boldsymbol{AS} \rangle_q \rceil \rangle_p$. Let $1 \leq i \leq d$ and $1 \leq j \leq \overline{n}$. If we denote the i-th row of $\boldsymbol{A}$ by $\boldsymbol{a}_i^T$ and the j-th column of $\boldsymbol{S}$ by $\boldsymbol{s}_j$, then

$$\boldsymbol{B}_{i,j} = \langle \lfloor \frac{p}{q} \langle \langle \boldsymbol{a}_i^T \boldsymbol{s}_j \rangle_q \rangle_q \rceil \rangle_p = \langle \lfloor \frac{p}{q} \langle \boldsymbol{a}_i^T \boldsymbol{s}_j \rangle_q \rceil \rangle_p.$$

By the definition of the rounding function $\lfloor \cdot \rceil$, we have that

$$\boldsymbol{B}_{i,j} \equiv \frac{p}{q} \langle \boldsymbol{a}_i^T \boldsymbol{s}_j \rangle_q + e_{i,j} \pmod{p} \text{ with } e_{i,j} \in (-1/2, 1/2].$$

As $\langle \boldsymbol{a}_i^T \boldsymbol{s}_j \rangle_q = \boldsymbol{a}_i^T \boldsymbol{s}_j + \lambda q$ for some integer λ, we infer that

$$\frac{q}{p} \boldsymbol{B}_{i,j} \equiv \boldsymbol{a}_i^T \boldsymbol{s}_j + \frac{q}{p} e_{i,j} \pmod{q}. \tag{10}$$

So we have d equations involving $\boldsymbol{s}_j$. Unlike conventional LWE, the errors $\frac{q}{p} e_{i,j}$ reside in a bounded interval, namely $(-\frac{q}{2p}, \frac{q}{2p}]$. In what follows, we will only consider the case that p divides q.

1.7.4 Primal Attack

In (10), we write $\boldsymbol{s}$ for $\boldsymbol{s}_j$, denote by $\boldsymbol{b}$ the vector of length m with j-th component $\frac{q}{p}\boldsymbol{B}_{i,j}$, and with $\boldsymbol{A}_m$ the matrix consisting of the m top rows of $\boldsymbol{A}$. We then have, for $\boldsymbol{e} \in (-\frac{q}{2p}, \frac{q}{2p}]^m$

$$\boldsymbol{b} \equiv \boldsymbol{A}_m \boldsymbol{s} + \boldsymbol{e} \pmod{q} \tag{11}$$

so that $\boldsymbol{v} = (\boldsymbol{s}^T, \boldsymbol{e}^T, 1)^T$ is in the lattice Λ defined as

$$\Lambda = \{\boldsymbol{x} \in \mathbb{Z}^{d+m+1} : (\boldsymbol{A}_m|\boldsymbol{I}_m|-\boldsymbol{b})\boldsymbol{x} = \boldsymbol{0} \pmod{q}\} \tag{12}$$

of dimension $d' = d + m + 1$ and volume q^m [25, 7]. The attacker then searches for a short vector in Λ which hopefully equals $\boldsymbol{v}$, thus enabling him to recover the secret $\boldsymbol{s}$.

Lattice Rescaling: The lattice vector $\boldsymbol{v} = (\boldsymbol{s}, \boldsymbol{e}, 1)$ is unbalanced in that $\|\boldsymbol{s}\| \ll \|\boldsymbol{e}\|$. For exploiting this fact, a rescaling technique originally due to [16], and analyzed further in [35] and [5] is applied. Multiplying the first d columns of Λ's basis (see Eq. 12) with an appropriate scaling factor ω yields the following weighted or rescaled lattice,

$$\Lambda_\omega = \{\boldsymbol{x} \in \mathbb{Z}^{d+m+1} : \left((\omega \cdot \boldsymbol{A}_m^T)|\boldsymbol{I}_m|-\boldsymbol{b}\right)\boldsymbol{x} = \boldsymbol{0} \pmod{\text{q}}\} \tag{13}$$

in which the attacker then searches for the shortest vector, that he hopes to be equal to $\boldsymbol{v}_\omega = (\omega \cdot \boldsymbol{s}^T, \boldsymbol{e}^T, 1)^T$. This search is typically done by using a lattice reduction algorithm to obtain a reduced basis of the lattice, the first vector of which will be the shortest of that basis due to a common heuristic. We explain later in this section how to choose an appropriate value for ω in order to maximize the chances of the attack's success.

If the quality of the lattice reduction is good enough, the reduced basis will contain $\boldsymbol{v}_\omega$. The attack success condition is as in [7] assuming that BKZ [33, 90] with block-size b is used as the lattice reduction algorithm. The vector $\boldsymbol{v}_\omega$ will be detected if its projection $\widetilde{\boldsymbol{v}}_b$ onto the vector space of the last b Gram-Schmidt vectors of Λ is shorter than the expected norm of the $(d'-b)^{th}$ Gram-Schmidt vector $\widetilde{\boldsymbol{b}}_{d'-b}$, where d' is the dimension of Λ [7, Sec. 6.3],[25]. The condition that must be satisfied for the primal attack to succeed is therefore:

$$
\begin{aligned}
&\quad \|\widetilde{\boldsymbol{v}}_b\| < \|\widetilde{\boldsymbol{b}}_{d'-b}\| \\
\text{i.e.,} \quad &\quad \|\widetilde{\boldsymbol{v}}_b\| < \delta^{2b-d'-1} \cdot (\text{Vol}(\Lambda))^{\frac{1}{d'}} \\
\text{where,} \quad &\quad \delta = \left((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi \mathbf{e}}\right)^{\frac{1}{2(b-1)}}
\end{aligned}
\tag{14}
$$

The attack success condition (14) yields the following *security* condition that must be satisfied by the parameters of our public-key encryption and key-

encapsulation schemes to remain secure against the primal attack:

$$\sqrt{(\omega^2 \cdot h + \sigma'^2 m) \cdot \frac{b}{d+m}} \geq \delta^{2b-d'-1} \cdot (q^m \omega^d)^{\frac{1}{d'}}$$

where,
$$\delta = ((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi \mathbf{e}})^{\frac{1}{2(b-1)}}, \qquad (15)$$

$$\sigma'^2 = ((q/p)^2 - 1)/12,$$

and
$$d' = d + m + 1.$$

For finding an appropriate value for ω, we rewrite (15) as

$$\delta^{2b-d'-1} b^{-1/2} \leq \sqrt{\omega^2 h + m\sigma'^2 \cdot \frac{1}{m+d}} \omega^{-d/d'} q^{-(d-d'-1)/d}. \qquad (16)$$

Given d, m, h and σ', the attacker obtains the least stringent condition on the block size b by maximizing the right hand side 16 over ω, or equivalently, by maximizing

$$\frac{1}{2} \log(\omega^2 h + m\sigma'^2) - \frac{d}{d'} \log \omega.$$

By differentiating with respect to ω, we find that the optimizing value for ω satisfies

$$\omega^2 = \frac{dm\sigma'^2}{h(d'-d)} = \frac{dm\sigma'^2}{h(m+1)} \approx \frac{d}{h}\sigma'^2.$$

1.7.5   Dual Attack

The dual attack against LWE/LWR [8],[6] employs a short vector $(\boldsymbol{v}, \boldsymbol{w}) \in \mathbb{Z}^m \times \mathbb{Z}^d$ in the dual lattice

$$\Lambda^* = \{(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{Z}^m \times \mathbb{Z}^d : \boldsymbol{A}_m^T \boldsymbol{x} = \boldsymbol{y} \;(\text{mod } q)\} \qquad (17)$$

The attacker constructs the distinguisher using $z = \{\boldsymbol{v}^T \boldsymbol{b}\}_q$. If $\boldsymbol{b}$ is an LWR samples, then $z = \{\boldsymbol{v}^T(\boldsymbol{A}_m \boldsymbol{s} + \boldsymbol{e})\}_q = \{\boldsymbol{w}^T \boldsymbol{s} + \boldsymbol{v}^T \boldsymbol{e}\}_q$. Note that $\|\boldsymbol{s}\| \ll \|\boldsymbol{e}\|$ in our case, the attacker can enforce that $\|\boldsymbol{w}\| \gg \|\boldsymbol{v}\|$ to ensure that $|\boldsymbol{w}^T \boldsymbol{s}| \approx |\boldsymbol{v}^T \boldsymbol{e}|$ similar to [6]. He does so by choosing $\omega = \sigma'\sqrt{m/h}$ (for the LWR rounding error with variance $\sigma'^2 = ((q/p)^2 - 1)/12$), and considering the lattice:

$$\Lambda_\omega^* = \{(\boldsymbol{x}, \boldsymbol{y}/\omega) \in \mathbb{Z}^m \times (\frac{1}{\omega} \cdot \mathbb{Z}^d) : \boldsymbol{A}_m^T \boldsymbol{x} = \boldsymbol{y} \;(\text{mod } q)\} \qquad (18)$$

A short vector $(\boldsymbol{v}, \boldsymbol{w}) \in \Lambda_\omega^*$ gives a short vector $(\boldsymbol{v}, \omega\boldsymbol{w}) \in \Lambda^*$ that is used to construct the distinguisher z. If $\boldsymbol{b}$ is uniform modulo q, so is z. If $\boldsymbol{b}$ is an LWR sample, then $z = \{\boldsymbol{v}^T \boldsymbol{b}\}_q = \{(\omega\boldsymbol{w})^T \boldsymbol{s} + \boldsymbol{v}^T \boldsymbol{e}\}_q = \{\boldsymbol{w}^T(\omega\boldsymbol{s}) + \boldsymbol{v}^T \boldsymbol{e}\}$ has a distribution approaching a Gaussian of zero mean and variance $\|(\boldsymbol{v}, \boldsymbol{w})\|^2 \cdot \sigma'^2$ as the lengths of the vectors increase, due to the Central limit theorem. Note that ω has been chosen such that $\|\omega\boldsymbol{s}\| \approx \|\boldsymbol{e}\|$. The maximal statistical distance between this and the uniform distribution modulo q is bounded by

$\epsilon \approx 2^{-1/2} \exp(-2\pi^2(\|(\boldsymbol{v}, \boldsymbol{w})\|^2 \cdot \sigma'/q)^2)$ [22, Appendix B]. Lattice reduction with root-Hermite factor δ yields a shortest vector of length $\delta^{d'-1} \cdot \text{Vol}(\Lambda_\omega^*)^{1/d'}$, where $d' = m + d$ and $\text{Vol}(\Lambda_\omega^*) = (q/\omega)^d$ are Λ_ω^*'s dimension and volume, respectively.

However, finding only one such short vector is not enough, as the resulting distinguishing advantage ϵ is too small to distinguish a final key which is hashed. The attack needs to be repeated approximately ϵ^{-2} times to achieve constant advantage. The attacker can utilize the fact that when sieving algorithms are used in the lattice reduction algorithm (e.g., BKZ with block size b) to obtain short vectors, each run of such algorithms provides $2^{0.2075b}$ vectors [8]. Under the conservative assumption that each of these vectors is short enough to guarantee a distinguishing advantage ϵ, the lattice reduction algorithm must therefore be run at least $\max(1, 1/2^{0.2075b} \cdot \epsilon^2)$ times [8], when considering BKZ with block size b. The cost of the weighted dual attack on LWR (with dimension d, large modulus q, rounding modulus p) using m samples thus is the cost of running one instance of BKZ lattice reduction with block size b times the number of repetitions required, i.e.,

$$(b \cdot 2^{cb}) \cdot \max(1, 1/(\epsilon^2 \cdot 2^{0.2075 \cdot b})), \text{ where}$$

$$\epsilon = 2^{-1/2} \cdot \mathbf{e}^{-2\pi^2 \left((\|(\boldsymbol{v}, \boldsymbol{w})\|^2 \cdot \sigma')/q \right)^2}, \quad \|(\boldsymbol{v}, \boldsymbol{w})\|^2 = \delta^{m+d-1} \cdot (q/\omega)^{d/(m+d)}, \quad (19)$$

$$\delta = ((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi \mathbf{e}})^{\frac{1}{2(b-1)}}, \quad \omega = \sigma' \cdot \sqrt{m/h}, \text{ and } \sigma'^2 = ((q/p)^2 - 1)/12.$$

The first term $(b \cdot 2^{cb})$ in the overall attack cost is that of running BKZ with block-size b. The BKZ sieving exponent c depends on the security model, e.g., $c = 0.265$ considering sieving algorithms sped up with Grover's quantum search algorithm [64, 67], and $c = 0.292$ with no such speedup.

For Round5's ring-based instantiations (i.e., when the parameters $n = d$), we follow a conservative approach similar to [8] and omit the b factor in the cost of running one instance of BKZ, accounting for the possibility that techniques in [88, 29] can be adapted to more advanced sieving algorithms, yielding the cost 2^{cb} instead of $b \cdot 2^{cb}$.

1.7.6 Hybrid Attack

In this section, we consider a hybrid lattice reduction and meet-in-the-middle attack (henceforth called *hybrid attack*) originally due to [57] that targeted the NTRU [54] cryptosystem. We first describe the hybrid attack, using notation similar to that of [94], in a general form. Subsequently, we specialize the attack to our scheme. Finally, we describe a scaling approach to make the hybrid attack exploit the fact that the secret is small and sparse.

The hybrid attack will be applied to the lattice

$$\Lambda' = \{\boldsymbol{x} \in \mathbb{Z}^{m+d+1} \mid (\boldsymbol{I}_m | \boldsymbol{A}_m | - \boldsymbol{b})\boldsymbol{x} \equiv 0 \bmod q\}$$

for some $m \in \{1, \ldots, d\}$. We first find a basis $\boldsymbol{B}'$ for Λ' of the form

$$\boldsymbol{B}' = \begin{pmatrix} \boldsymbol{B} & \boldsymbol{C} \\ \boldsymbol{0} & \boldsymbol{I}_r \end{pmatrix} \tag{20}$$

where $0 < r < d$ is the meet-in-the-middle dimension and $\boldsymbol{I}_r$ is the r-dimensional identity matrix. We aim to find the short vector $\boldsymbol{v} = (\boldsymbol{e}^T, \boldsymbol{s}^T, 1)^T$ in Λ'. We write $\boldsymbol{v} = (\boldsymbol{v}_l^T \boldsymbol{v}_g^T)^T$ where $\boldsymbol{v}_g$ has length r. We recover the vector $\boldsymbol{v}_g$ of length $r < d$ consisting of the $(r-1)$ bottom entries of $\boldsymbol{s}$ followed by a '1' by guessing. As the columns of $\boldsymbol{B}'$ generate Λ', there exists a $\boldsymbol{x} \in \mathbb{Z}^{d+m+1-r}$ such that

$$\boldsymbol{v} = \begin{pmatrix} \boldsymbol{v}_l \\ \boldsymbol{v}_g \end{pmatrix} = \boldsymbol{B}' \begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{v}_g \end{pmatrix} = \begin{pmatrix} \boldsymbol{B} & \boldsymbol{C} \\ \boldsymbol{0} & \boldsymbol{I}_r \end{pmatrix} \begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{v}_g \end{pmatrix} = \begin{pmatrix} \boldsymbol{Bx} + \boldsymbol{Cv}_g \\ \boldsymbol{v}_g \end{pmatrix} \tag{21}$$

As $\boldsymbol{v}_l$ is short, $\boldsymbol{Cv}_g$ is close to $-\boldsymbol{Bx}$, a vector from the lattice $\Lambda(\boldsymbol{B})$. As explained in [94], the idea is that if we correctly guess $\boldsymbol{v}_g$, we can hope to find $\boldsymbol{v}_l$ by using Babai's Nearest Plane (NP) algorithm [15]. This algorithm, given a basis $\tilde{\boldsymbol{B}}$, finds for every target vector $\boldsymbol{t} \in \mathbb{R}^{d+m+1-r}$ a vector $\boldsymbol{e} = \mathrm{NP}_{\tilde{\boldsymbol{B}}}(\boldsymbol{t})$ such that $\boldsymbol{t} - \boldsymbol{e} \in \Lambda(\tilde{\boldsymbol{B}})$.

The cost for the hybrid attack thus is the sum of two terms: the cost of finding a good basis $\tilde{\boldsymbol{B}}$ for $\Lambda(\boldsymbol{B})$, and the cost of generating $\mathrm{NP}_{\tilde{\boldsymbol{B}}}(\boldsymbol{Cy})$ for all $\boldsymbol{y}$ from a set of vectors of length r that (with high probability) contains $\boldsymbol{v}_g$. The latter cost may be reduced by using a Meet-in-the-middle approach [57] which reduces the number of calls to the Nearest Plane algorithm to the square root of the number of calls with a brute-force approach.

As $r < d$, the vector $\boldsymbol{v}_g$ is a ternary. The attacker can benefit from the fact that $\boldsymbol{s}$ has h non-zero entries in the generation of candidates for $\boldsymbol{v}_g$: candidates with high Hamming weight are not very likely. Also, as the $(d-r)$ bottom entries of $\boldsymbol{v}_l$, being the $(d-r)$ top elements of $\boldsymbol{s}$, are ternary and sparse. In order to benefit from this fact, the $d-r$ rightmost columns of the matrix $\boldsymbol{B}$ are multiplied with an appropriate scaling factor ω. Calculated similarly to Section 1.7.4 by equalizing the *per-component* expected norms of the secret $\boldsymbol{s}$ and LWR rounding error $\boldsymbol{e}$, we arrive at the same scaling factor $\omega = \sqrt{\frac{(q/p)^2-1}{12} \cdot \frac{d}{h}}$. This then scales up the volume of the $(d - r + m + 1$ dimensional) lattice Λ generated by $\boldsymbol{B}$ by a factor ω^{d-r}.

We analyze the hybrid attack similarly as in [54]. For each pair (r, m) with $1 \leq r, m < d$, we stipulate that the quality of the reduced basis $\tilde{\boldsymbol{B}}$ is so high that $\mathrm{NP}_{\tilde{\boldsymbol{B}}}(\boldsymbol{v}_g) = \boldsymbol{v}_l$ with high probability. The condition, derived from [57, Lemma 1] is that the norm of the last Gram-Schmidt vector of $\tilde{\boldsymbol{B}}$ is at least twice $\|\boldsymbol{v}_l\|_\infty$, see also [54] We use the Geometric Series Assumption to approximate the norms of these vectors in terms of the Hermite constant δ. The cost for obtaining a reduced basis with Hermite constant δ is estimated as $b2^{cb}$, where b is such that $\delta = \left((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi e} \right)^{\frac{1}{2b-1}}$, and similar to our conservative approach in the previous attacks we ignore the additional b factor in the cost accounting for

potential speedups in sieving algorithms for ring-based instantiations. c is the BKZ sieving exponent as before. Moreover, we estimate the cost for the lattice decoding part to be equal to the number of invocations of the Nearest Plane Algorithm, which, following [54], we set to $2^{\frac{1}{2}r \cdot H}$, where H is the entropy of the distribution of each of the coordinates of the guessed vector $\boldsymbol{v}_g$. The number $2^{r \cdot H}$ approximates the number of typical vectors of length r; the factor $\frac{1}{2}$ is due to either the usage of the MITM technique, or the use of Grover's algorithm in the quantum case. Finally, we minimize the cost over all feasible pairs (r, m).

Our analysis of the hybrid attack allows us to obtain a rough estimate of its cost. With the goal of providing a more accurate estimate, Wunderer [94] gives an extensive runtime analysis of the hybrid attack. One of the aspects he takes into account is that the attacker chooses a larger value of δ. This leads to a smaller cost (running time) for lattice reduction to obtain $\tilde{\boldsymbol{B}}$, but decreases the probability that $\text{NP}_{\tilde{\boldsymbol{B}}}(\boldsymbol{v}_g) = \boldsymbol{v}_l$, thereby increasing the expected cost (running time) of the part of the attack dealing with solving BDD problems. Also, he takes into account that the guesses for $\boldsymbol{v}_g$ are generated such that the most likely candidates for $\boldsymbol{v}_g$ occur early, thus reducing the expected number of calls to the nearest plane algorithm.

1.7.7 Attacks against Sparse Secrets

In Sections 1.7.4 and 1.7.5, we considered attacks [16, 5, 35] against LWE and/or LWR variants with unnaturally small secrets. In this section, we consider attacks against *sparse* secrets with the goal of choosing an appropriate Hamming weight providing both optimal performance and security. The best-known attacks against such sparse secrets are (to the best of our knowledge) the Hybrid attack described in Section 1.7.6, and another one due to Albrecht *et al* [5]. The Hybrid attack performs better than Albrecht's attack against our schemes, it is therefore the primary attack considered in our analysis. Recall that the hybrid attack's overall cost is the sum of two components: that of finding a good basis and that of solving Babai's Nearest Planes for a large set of vectors using a Meet-in-the-middle approach. Recall also that this second cost component depends on the entropy H of the distribution of each secret coordinate (in our case), which in turn depends on the Hamming weight of the secret. We therefore optimize over the Hamming weight to choose the smallest value for which the overall hybrid attack cost is at least a targeted minimum (depending on the security level).

For completeness, we also describe Albrecht's attack [5] against LWE/LWR variants with sparse secrets. Since most rows of the public matrix $\boldsymbol{A}$ become irrelevant in the calculation of the product $\boldsymbol{As}$ for such secrets, Albrecht's attack ignores a random set of $k \leq d$ components of $\boldsymbol{s}$ and brings down the lattice dimension (and hence attack cost) during lattice reduction. As $\boldsymbol{s}$ has $d - h$ zeroes, there are $\binom{d-h}{k}$ choices (out of $\binom{d}{k}$) for the k ignored components such that these ignored components only contain zeroes. We therefore repeat the attack $\binom{d}{k} / \binom{d-h}{k}$ times. We estimate the cost (in bits) for a given Hamming weight $h \leq d$, as the number of repetitions each low cost attack is performed

times the cost of the low-cost attack on a lattice of dimension $d - k$:

$$\frac{\binom{d}{k}}{\binom{d-h}{k}} \times \text{Cost}_{\text{Lattice Reduction}}(d, k, h)$$

Here $\text{Cost}_{\text{Lattice Reduction}}(d, k, h)$ is defined as

$$\min\{b \cdot 2^{cb} \mid b \in \mathbb{N}, \text{ there exists } m \in \mathbb{N} \text{ such that } m \leq d \text{ and (22) is satisfied.}\}$$

$$\sqrt{(\omega^2 \cdot h + \sigma'^2 m) \cdot \frac{b}{d+m}} < \delta^{2b-d'-1} \cdot (q^{d'-(d-k)-1}\omega^{d-k})^{\frac{1}{d'}}$$

$$\text{where,} \qquad \delta = ((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi \mathbf{e}})^{\frac{1}{2(b-1)}},$$

$$\omega = \sigma' \cdot \sqrt{(d-k)/h},$$

$$\sigma'^2 = ((q/p)^2 - 1)/12,$$

$$\text{and} \qquad d' = m + d - k + 1.$$

(22)

The term $b \cdot 2^{cb}$ represents the cost of running BKZ lattice reduction with block-size b (the first term or factor b in the cost is ignored for ring-based instantiations as in the previous attacks). The attack runs on a LWE problem of dimension $d - k$ with $m \leq d$ samples. Condition 22, which is essentially Condition 14, ensures that such an attack has chances of succeeding. Note that although the above applies to the primal attack, a similar analysis is possible for the dual attack, in which case $\text{Cost}_{\text{Lattice Reduction}}(d, k, h)$ is calculated as in Eq. 19, with the parameter d replaced by $d - k$.

This specialized attack only gives an advantage if an attacker is able to choose a k for which the total attack cost is less than a standard lattice-reduction attack on a lattice of dimension d. Similar to the case of the hybrid attack (Section 1.7.6), we optimize over the Hamming weight to choose the smallest value such that Albrecht et al.'s attack results in at least a minimum targeted security level (both for the standard attack embodiment mentioned above and an adaptive embodiment described in [5]).

Furthermore, to ensure that an exhaustive, brute-force search of each secret-key vector in the secret-key using Grover's quantum search algorithm [49] has a cost of at least λ (in bits), the chosen Hamming weight should satisfy:

$$\sqrt{\binom{d}{h} \cdot 2^h} > 2^\lambda$$

(23)

Note that for a typical security level of $\lambda = 128$, a dimension of at least $d = 512$ would be secure against Grover's quantum search, for any Hamming weight h that is at least $0.1d$.

1.7.8 Pre-computation and Back-door Attacks

A pre-computation attack can happen if the GLWR public parameter $\boldsymbol{A}$ is fixed and an attacker performs lattice reduction on it over a long period of time. A

back-door attack can happen if there is any public value, e.g., $\boldsymbol{A}$ is deliberately chosen so as to result in a lattice with weak security. For each $\tau \in \{0, 1, 2\}$, the function $f_{d,n}^{(\tau)}$ (see Section 1.4.2) prevents both types of attacks.

For $\tau = 0$, a fresh $\boldsymbol{A}$ is generated in each protocol instantiation, similar to [25] and [7]. This prevents both pre-computation attacks and backdoors.

For $\tau = 1$, the matrix $\boldsymbol{A}$ is derived from a fixed long-term matrix $\boldsymbol{A}_{\text{master}}$ of dimension d^2. This is done by applying a random, fresh permutation that is chosen by the initiator of the protocol at the start of each protocol exchange. This prevents any pre-computation attacks since the possible number of per-muted $\boldsymbol{A}$ obtained in this way equals n^n. Back-doors are avoided since $\boldsymbol{A}_{\text{master}}$ is derived by means of a pseudo-random function from a randomly generated seed. We note that for both $f_{d,n=1}^{(0)}$ and $f_{d=n=1}^{(1)}$, the entries of both $\boldsymbol{A}_{\text{master}}$ and the resulting $\boldsymbol{A}$ cannot be differentiated from uniform, hence the results in Section A.8 hold.

For $\tau = 2$, $\boldsymbol{A}$ is obtained from a vector $\boldsymbol{a}_{\text{master}}$ of length q, that is randomly generated by means of a DBRG from a seed determined by the initiator in each protocol interaction. Furthermore, each row in $\boldsymbol{A}$ is obtained from this vector by means of a random permutation that is also determined by the ini-tiator and is specific to each protocol interaction. Since only a few elements need to be generated and kept in memory, $f_{d,n}^{(2)}$ is efficient. Pre-computation and back-door attacks are avoided since the seed that determines $\boldsymbol{A}$ is new in each protocol interaction. Furthermore, this approach destroys any struc-ture in the resulting $\boldsymbol{A}$ (as can be found in circulant or anti-circulant matrices for ideal lattices, for example) since it contains many more elements. This results clear from the following analysis. Suppose $a = \langle a_0, a_1, a_2, \ldots, a_{n-1} \rangle$ and $b = \langle b_0, b_1, \ldots, b_{k-1}, a_0, a_1, a_2, \ldots, a_{n-k-1} \rangle$ are two rows of $\boldsymbol{A}$. They share $n - 1 - k$ elements due to our rotation strategy. Then, define $a(x) := a_0 + a_1 x + a_2 x^2 + \cdots + a_{n-1} x^{n-1}$ and $b(x) := b_0 + b_1 x + \cdots + b_{k-1} x^{k-1} + a_0 x^k + a_1 x^{k+1} + \cdots + a_{n-k-1} x^{n-1}$. We have $b(x) = a(x) x^k + (a_{n-k} + b_0) + (a_{n-k+1} + b_1) x + \cdots + (a_{n-1} + b_{k-1}) x^{k-1} \mod x^n + 1$. Effectively, that is, each row can be seen as using the $x^n + 1$ ring with a random shifting (due to k) and additional random noises, those $(a_{n-k} + b_0) + (a_{n-k+1} + b_1) x + \cdots + (a_{n-1} + b_{k-1}) x^{k-1}$. From this point of view, generating $\boldsymbol{A}$ with $f_{d,n}^{(2)}$ is at least as secure as using the pure ring version of Round5. Since the length of $\boldsymbol{a}_{\text{master}}$ is smaller than d^2, the entries of the resulting $\boldsymbol{A}$ can be differentiated from uniform. The security of Round5 can therefore not be based on the results in Section A.8 for $f_{d,n}^{(2)}$.

1.8   Correctness of Round5

The following subsections we first derive the condition for correct decryption in Round5. Subsequently, we work out this condition for non-ring parameters, ring parameters with reduction polynomial $\xi(x) = \Phi_{n+1}(x)$, and ring parame-ters with $\xi(x) = x^{n+1} - 1$. The analysis for ring parameters with $\xi(x) = \Phi_{n+1}(x)$ show that decryption errors (before error correction) are correlated. For this rea-son, error correction in combination with reduction polynomials $\xi(x) = x^{n+1} - 1$

is not effective. Finally, results from the decryption failure analysis are compared with simulations for scaled-down versions of Round5.

1.8.1 Decryption failure analysis

In this section, the decryption failure behavior of r5_cpa_pke is analyzed. In decryption, the vector $\boldsymbol{x'} = \langle \lfloor \frac{b}{p} \boldsymbol{\zeta} \rceil \rangle_b$ is computed, where

$$\boldsymbol{\zeta} = \left\langle \frac{p}{t}\boldsymbol{v} + h_3\boldsymbol{j} - \text{Sample}_\mu \left(\left\langle \boldsymbol{S}^T(\boldsymbol{U} + h_4\boldsymbol{J}) \right\rangle_\xi \right) \right\rangle_p.$$

First, a sufficient condition is derived so that $\boldsymbol{x'}$ and $\boldsymbol{x} = xef\_compute_{\kappa,f}(m)$ agree in a given position, where $\boldsymbol{x}$ is considered as a vector of μ symbols, each consisting of b_bits bits. We have that

$$\boldsymbol{v} \equiv \langle \frac{t}{p}\text{Sample}_\mu(\langle \boldsymbol{B}^T \boldsymbol{R} \rangle_\xi + h_2\boldsymbol{j}) - \frac{t}{p}\boldsymbol{I}_v \rangle_p + \frac{t}{b}\boldsymbol{x} \pmod{t},$$

where $\frac{t}{p}\boldsymbol{I}_v$ is the error introduced by the rounding downwards, with each component of $\boldsymbol{I}_v$ in $\mathbb{Z}_{p/t}$. As a result,

$$\boldsymbol{\zeta} \equiv \frac{p}{b}\boldsymbol{x} + \boldsymbol{\Delta} \pmod{p} \text{ with } \boldsymbol{\Delta} = (h_2+h_3)\boldsymbol{j} - \boldsymbol{I}_v + \text{Sample}_\mu(\langle \boldsymbol{B}^T\boldsymbol{R} - \boldsymbol{S}^T(\boldsymbol{U}+h_4\boldsymbol{J}\rangle_\xi). \tag{24}$$

As $\boldsymbol{x'} = \lfloor \frac{b}{p}\boldsymbol{\zeta} \rceil = \lfloor \frac{b}{p}\boldsymbol{\zeta} - \frac{1}{2} \rceil$, we have that

$$\boldsymbol{x'} \equiv \boldsymbol{x} + \lfloor \frac{b}{p}\boldsymbol{\Delta} - \frac{1}{2}\boldsymbol{J} \rceil \equiv \boldsymbol{x} + \lfloor \frac{b}{p}\{\boldsymbol{\Delta} - \frac{p}{2b}\boldsymbol{J}\}_p \rceil \pmod{b}.$$

Here $\{y\}_p$ denotes the integer in $(-p/2, p/2]$ that is equivalent to y modulo p. As a consequence, $x_i = x'_i$ whenever $|\{\boldsymbol{\Delta}_i - \frac{p}{2b}\}_p| < \frac{p}{2b}$. We multiply both sides of the above inequality with $\frac{q}{p}$, and infer that $x_i = x'_i$ whenever

$$| \{\frac{q}{p}\boldsymbol{\Delta}_i - \frac{q}{2b}\}_q | < \frac{q}{2b}. \tag{25}$$

Equivalently, as $\frac{q}{p}\boldsymbol{\Delta}_i$ has integer components, if $x_i \neq x'_i$, then

$$\langle \frac{q}{p}\boldsymbol{\Delta}_i \rangle_q \in \left[\frac{q}{2b}, q - \frac{q}{2b} \right] \tag{26}$$

The probability that $\boldsymbol{x}$ and $\boldsymbol{x'}$ differ in position i thus is at most the probability that (26) is satisfied. In order to analyze this probability, we work out $\frac{q}{p}\boldsymbol{\Delta} - \frac{q}{2b}\boldsymbol{J}$, using (24).

We write $\boldsymbol{J}_v = \frac{q}{p}(h_2\boldsymbol{j} + h_3\boldsymbol{j} - \boldsymbol{I}_v - \frac{p}{2^b}\boldsymbol{J})$. The definitions of h_2 and h_3 imply that $\boldsymbol{J}_v = \frac{q}{p}(\frac{p}{2t} - \boldsymbol{I}_v)$. Each component of $\boldsymbol{I}_v$ is in $\mathbb{Z}_{p/t}$. The value of h_3 thus ensures that the absolute value of each coefficient of $\frac{p}{2t} - \boldsymbol{I}_v$ is at most $\frac{p}{2t}$.

We now analyse $\frac{q}{p}\langle \boldsymbol{B}^T\boldsymbol{R} - \boldsymbol{S}^T(\boldsymbol{U}+h_4\boldsymbol{J})\rangle_\xi$. Similarly to the expression for $\boldsymbol{v}$, we write

$$\boldsymbol{B} = \left\langle \frac{p}{q}\left(\langle \boldsymbol{AS}\rangle_{\Phi_{n+1}} + h_1\boldsymbol{J}\right) - \frac{p}{q}\boldsymbol{I}_B \right\rangle_p \text{ and } \boldsymbol{U} = \left\langle \frac{p}{q}(\langle \boldsymbol{A}^T\boldsymbol{R}\rangle_{\Phi_{n+1}} + h_2\boldsymbol{J}) - \frac{p}{q}\boldsymbol{I}_U \right\rangle_p,$$

with all components of $\boldsymbol{I}_B$ and $\boldsymbol{I}_U$ in $\mathbb{Z}_{q/p}$. We thus can write

$$\frac{q}{p}(\boldsymbol{B}^T\boldsymbol{R} - \boldsymbol{S}^T(\boldsymbol{U} + h_4\boldsymbol{J}) \equiv \langle \boldsymbol{S}^T\boldsymbol{A}^T\rangle_{\Phi_{n+1}}\boldsymbol{R} - \boldsymbol{S}^T\langle \boldsymbol{A}^T\boldsymbol{R}\rangle_{\Phi_{n+1}} + \boldsymbol{J}_B^T\boldsymbol{R} - \boldsymbol{S}^T\boldsymbol{J}_U \quad (\text{mod } q) \tag{27}$$

where

$$\boldsymbol{J_B} = h_1\boldsymbol{J} - \boldsymbol{I}_B, \text{ and } \boldsymbol{J_U} = (h_2 + h_4)\boldsymbol{J} - \boldsymbol{I}_U. \tag{28}$$

As $h_1 = h_2 + h_4 = \frac{q}{2p}$, all entries of $\boldsymbol{J}_B$ and of $\boldsymbol{J}_U$ are from the set $I :=$ $(-\frac{q}{2p}, \frac{q}{2p}] \cap \mathbb{Z}$. The value of h_4 thus ensures that the absolute value of the entries of $\boldsymbol{J}_B$ and $\boldsymbol{J}_U$ are at most $\frac{q}{2p}$.

Clearly, if $\xi = \Phi_{n+1}$, then $\langle \boldsymbol{S}^T\boldsymbol{A}^T\rangle_{\Phi_{n+1}}\boldsymbol{R} \equiv \boldsymbol{S}^T\langle \boldsymbol{A}^T\boldsymbol{R}\rangle_{\Phi_{n+1}}$, so that

$$\frac{q}{p}\boldsymbol{\Delta} \equiv \boldsymbol{J}_v + \text{Sample}_\mu\left(\langle \boldsymbol{J}_B^T\boldsymbol{R} - \boldsymbol{S}^T\boldsymbol{J}_U\rangle_{\Phi_{n+1}}\right) \quad (\text{mod } q) \tag{29}$$

In Section 1.8.4, the special case $\xi(x) = x^{n+1} - 1$ will be analyzed.

The probability that (26) is satisfied will be analyzed under the following assumptions: the entries $\boldsymbol{J}_v$ are drawn independently, and each such entry is distributed as $\frac{q}{p}y$ with y uniform on $(-\frac{p}{2t}, \frac{p}{2t}] \cap \mathbb{Z}$. The entries of $\boldsymbol{J}_B$ and $\boldsymbol{J}_U$ are drawn independently and uniformly from $I = (-\frac{q}{2p}, \frac{q}{2p}] \cap \mathbb{Z}$. In our analysis, we also assume that all columns of the secret matrices $\boldsymbol{S}$ and $\boldsymbol{R}$ have $h/2$ coefficients equal to 1 and $h/2$ coefficients equal to -1. This is true for the implementations of f_S and f_R.

1.8.2 Failure probability computation: non-ring parameters

In the non-ring case, each entry of $\boldsymbol{J}_B^T\boldsymbol{R}$ and of $\boldsymbol{S}^T\boldsymbol{J}_U$ is the inner product of a row of $\boldsymbol{J}_B^T$ (resp. a column of $\boldsymbol{J}_U$) and a ternary vector with $h/2$ entries equal to one and $h/2$ entries equal to minus one. Hence each entry of $\boldsymbol{J}_B^T\boldsymbol{R} - \boldsymbol{S}^T\boldsymbol{J}_U$ is distributed as the sum of h uniform variables on I minus the sum of h uniform variables on I. Assuming independence, the latter distribution (modulo q) can easily be computed explicitly (using repeated convolutions). Indeed, if c and d are two independent variables on $\{0, 1, \ldots, q-1\}$ with probability distributions p_c and p_d, then the probability distribution p_e of $\langle c + d\rangle_q$ is given by

$$p_e(k) = \sum_{i=0}^{q-1} p_c(i)p_d(\langle k - i\rangle_q) \text{ for } 0 \le k \le q - 1.$$

Assuming independence between the i-th components of $\boldsymbol{J}_v$ and of $\text{Sample}_\mu(\boldsymbol{J}_B^T\boldsymbol{R} - \boldsymbol{S}^T\boldsymbol{J}_U)$, the probability that (26) is satisfied can be computed by using another convolution. By the union bound, the probability that at least one of the symbols of $\boldsymbol{x}$ is not retrieved correctly is at most μ times the probability that (26) is satisfied.

1.8.3 Failure probability computation: ring parameters with $\xi(x) = x^{n+1} - 1$

Round5 parameters without error correction use $\xi(x) = \Phi_{n+1}(x)$. As $N(x)$ is a multiple of $\Phi_{n+1}(x)$, we have that $\langle f \rangle_{\Phi_{n+1}} = \langle \langle f \rangle_N \rangle_{\Phi_{n+1}}$. Moreover, if $g(x) = \sum_{i=0}^{n} g_i x^i$, then $\langle g \rangle_{\Phi_{n+1}} = g - g_n \Phi_{n+1}$. In particular, for all polynomials s, e,

$$\text{if } \langle se \rangle_N = \sum_{k=0}^{n} c_k(s,e)x^k, \text{ then } \langle se \rangle_{\Phi_{n+1}} = \sum_{k=0}^{n-1} (c_k(s,e) - c_n(s,e))x^k, \quad (30)$$

with $c_k(s,e)$ as defined in (33). If the k-th symbol is not retrieved correctly, then

$$\langle (j_v(x))_k + c_k(j_B, r) - c_n(j_B, r) + c_k(s, j_u) + c_n(s, j_u) \rangle_q \in \left[\frac{q}{2b}, q - \frac{q}{2b} \right] \quad (31)$$

Assuming independence, and taking into account that r and s contain $h/2$ ones and $h/2$ minus ones, $c_k(j_B, r) - c_n(j_B, r) - c_n(j_B, r) + c_k(s, j_u)$ is distributed as the difference of $2h$ independent random variables on I, minus the sum of $2h$ independent random variables on I. The probability that (31) is satisfied thus can be computed explicitly. By the union bound, the probability that at least one of the μ symbols is not retrieved correctly is at most μ times the probability that (31) is satisfied.

Remark The condition in (31) for decryption error in position k shows the term $-c_n(j_B, r) + c_n(s, j_u)$ that is common to all positions k. This is the reason that using error correction in conjunction with $\xi(x) = \Phi_{n+1}(x)$ is not effective. The union bound can be used as it also applies to dependent random variables.

1.8.4 Failure probability computation: ring parameters with $\xi(x) = \Phi_{n+1}(x)$

Round5 parameters using error correction are restricted to the ring case. So in (27) we have

$$\langle \boldsymbol{S}^T \boldsymbol{A}^T \rangle_{\Phi_{n+1}} \boldsymbol{R} - \boldsymbol{S}^T \langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}} = \lambda_s(x)\Phi_{n+1}(x)r(x) - s(x)\lambda_r(x)\Phi_{n+1}(x)$$

for some $\lambda_s, \lambda_r \in \mathbb{Z}[x]$.

For the Round5 parameters with error correction, it is required that $\xi(x) = (x-1)\Phi_{n+1}(x) = x^{n+1} - 1$, and that $s(x)$ and $r(x)$ both are balanced, that is, have $h/2$ coefficients equal to 1, $h/2$ coefficients equal to -1, and the other coefficients equal to zero. Then $x - 1$ divides both $r(x)$ and $s(x)$, and so $\xi(x)$ divides both $\Phi_{n+1}(x)r(x)$ and $s(x)\Phi_{n+1}(x)$. As a result, (29) reads as

$$\frac{q}{p}\Delta(x) \equiv j_v(x) + \text{Sample}_\mu \left(\langle \boldsymbol{j}_B(x)r(x) - s(x)j_U(x) \rangle_{x^{n+1}-1} \right) \pmod{q}. \quad (32)$$

For any two polynomials $e(x) = \sum_{i=0}^{n} e_i x^i$ and $s(x) = \sum_{i=0}^{n} s_i x^i$,

$$\langle s(x)e(x)\rangle_{x^{n+1}-1} = \sum_{k=0}^{n} c_k(s,e)x^k \text{ where } c_k(s,e) = \sum_{i=0}^{n} e_i s_{\langle k-i\rangle_{n+1}}. \qquad (33)$$

If the k-th symbol is not retrieved correctly, then

$$\langle (j_v(x))_k + c_k(j_B, r) - c_k(s, j_u)\rangle_q \in \left[\frac{q}{2b}, q - \frac{q}{2b}\right]. \qquad (34)$$

Assuming independence, using that r and s both contain $h/2$ and $h/2$ minus ones, $c_k(j_B, r) - c_k(s, j_u)$ is distributed as the sum of h independent uniform variables on I, minus the sum of h uniform variables on I. The probability p_n that (33) is satisfied thus can be computed explicitly.

Now let the error-correcting code be capable of correcting f symbol errors. Assuming that $c_k(s,e)$ and $c_\ell(s,e)$ are independent whenever $k \neq \ell$, the probability of not decoding correctly is at most $\sum_{e \geq f+1} \binom{\mu}{e} p_n^e (1 - p_n)^{\mu-e}$.

1.8.5 Experimental results

Figure 3 compares the estimated probabilities of at least one error occurring and that of at least two errors occurring, when $\xi = N_{n+1}$ (as in r5_cpa_pke) and when $\xi = \Phi_{n+1}$,respectively. These estimates are computed by explicitly convolving probability distributions. Parameters are simulated *without* error correction, and are $n = 800$, $q = 2^{11}$, $p = 2^7$, $t = 2^4$, $\mu = \kappa = 128$, while the Hamming weight h varies between 100 and 750 in order to show its effect on both the bit failure rate and error correlation.

For $xi = \Phi_{n+1}$, the probabilities are computed as follows. For any a, (31) can be used to compute $p(k \mid a)$, the probability that bit k is not retrieved correctly, given that $-c_n(j_b, r) + c_n(s, j_u) \equiv a \pmod{q}$. We assume that having a bit error in position k, given that $c_n(s, j_u) - c_n(j_b, r) \equiv a$, is independent of having a bit error in another position j, given that $c_n(s, j_u) - c_n(j_{b,r}) \equiv a$. The probability of having exactly e bit errors, given that $c_n(s, j_u) - c_n(j_b, r) \equiv a$, then equals $\binom{\mu}{k}(p(0 \mid a)^e(1 - p(0 \mid a))^{\mu-e}$. By summing these probabilities over a, weighted with the probability that $c_n(s, j_u) - c_n(j_b, r) \equiv a$, the probability of having exactly e bit errors is obtained.

Clearly, the probability of at least two errors is much higher when multiplications are done modulo Φ_{n+1} instead of N_{n+1}, and in the latter case, this probability is significantly lower than the probability of at least one error. Figure 3 also shows corresponding probabilities of at least one and at least two errors, obtained from simulations of *actual, scaled-down* r5_cpa_pke parameters, showing that the actual behavior closely matches estimates.

To conclude, the effect of dependency introduced between polynomial coefficients in Round2 due to polynomial multiplication modulo Φ_{n+1}, is made negligible by the combined use of polynomial multiplication modulo N_{n+1} and balanced secrets in Round5, allowing the use of forward error correction, resulting in better security and performance.
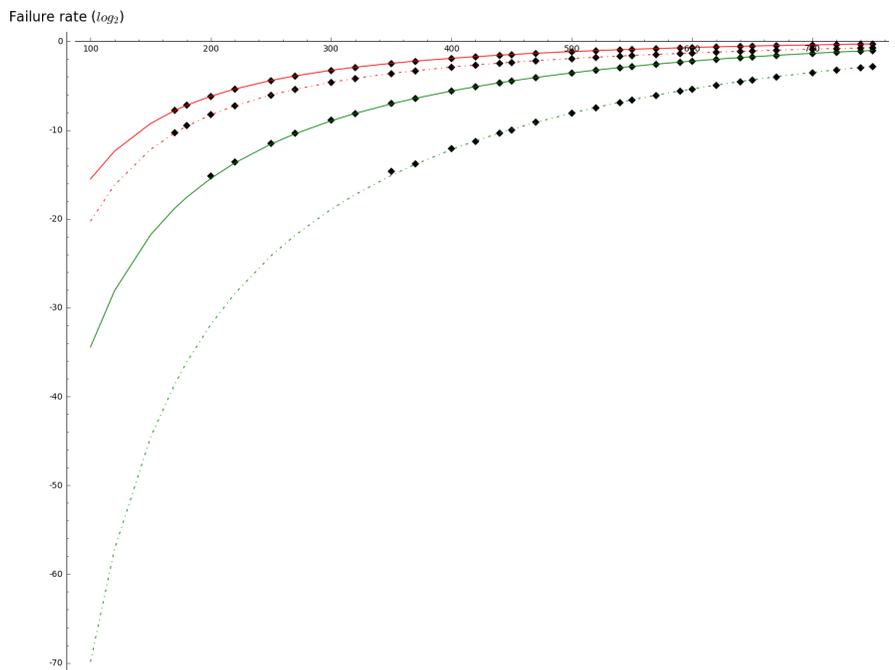
Figure 3: Probabilities of at least one (continuous lines) and at least two errors (dotted lines) in Round5 ring parameters, plotted against the Hamming weight of secrets (X-axis). Red and green curves respectively represent predicted values for the above in parameters using $\Phi_{n+1}(x)$ and $N_{n+1}(x)$ as reduction polynomial, respectively. Diamonds represent corresponding probabilities computed from actual Round5 simulations for the same parameters, with 10^6 runs per datapoint. Scripts for analyzing and reproducing these results can be found at www.round5.org.

1.9 Round5: Configurations, Parameters and Performance

Round5 offers a large design space allowing for a unified and efficient implementation of Round5, independently of the fact that it instantiates LWR or RLWR. r5_cpa_kem and r5_cca_pke can be configured to instantiate the LWR and RLWR underlying problems depending on the input configuration parameter n. As described in detail later, the security level depends on the input value d and also the choice of q and p. All moduli, q, p, and t are chosen to be powers-of-two. Usage of error correction to better deal with decryption failures can also be configured by setting up parameter f. If $f > 0$, then secrets must be balanced and the reduction polynomial $\xi(x)$ for the ciphertext component v must be $x^{n+1} - 1$. A final restriction is that $\Phi_{n+1}(x)$ is irreducible modulo two.

The Round5 configurations are described in Sections 1.9.1 and Section 1.9.2. Section 1.9.3 describes the implementations in the submission package. Sec-

tion 1.9.5 details the configuration parameters for Round5.KEM. Section 1.9.6 summarizes the performance of Round5.KEM. Section 1.9.7 details the configuration parameters of Round5.PKE. Section 1.9.8 summarizes the performance of Round5.PKE. Section **??** reports performance on other platforms.

1.9.1   Main Configurations of Round5

The Round5 cryptosystem includes a total of 18 main parameters sets: six ring configurations without error correction, six ring configurations without error correction, and six non-ring configurations. Each set of six configurations is for Round5.KEM and Round5.PKE, NIST security levels 1, 3 and 5.

A parameter set is denoted as R5N{1,D}_{1,3,5}{KEM,PKE}_{0,5}{version} where:

- 1,D refers whether it is a non-ring (1) or ring (D) parameter set.

- 1,3,5 refers to the NIST security level that is strictly fulfilled.

- KEM,PKE refers to the cryptographic algorithm it instantiates.

- 0,1,2,3,4,5 identifies the amount of corrected bits, 0 means no-errors are corrected and this description is equivalent to the original Round2; 5 means that up to 5 errors can be corrected as in the original HILA5 submission.

- version is a letter to indicate the version of published parameters to account for potential differences in published parameters.

R5ND_{1,3,5}KEM_5c: Merged parameters for the ring variant ($n = d$) of Round5 key-encapsulation mechanism, for NIST security levels 1, 3 and 5, resulting from the merge of the NIST PQC first round candidate cryptosystems Round2 and HILA5. XE5 forward error correction is used to decrease decryption failure rates (hence the 5 at the end of the parameter designator), and improve bandwidth and security.

R5ND_{1,3,5}PKE_5c: Merged parameters for the ring variant ($n = d$) of Round5 public-key encryption, for NIST security levels 1, 3 and 5, resulting from the merge of the NIST PQC first round candidate cryptosystems Round2 and HILA5. XE5 forward error correction is used to decrease decryption failure rates (hence the 5 at the end of the parameter designator), and improve bandwidth and security.

R5ND_{1,3,5}KEM_0c: Parameters for the ring variant ($n = d$) of Round5 INDCPA key-encapsulation mechanism, for NIST security levels 1, 3 and 5. No forward error correction is used, hence the 0 at the end of the parameter designator. All polynomial multiplications are done modulo the prime-order cyclotomic polynomial. This choice is considered more conservative since it uses the same design principles as Round2.

R5ND_{1,3,5}PKE_0c: Parameters for the ring variant ($n = d$) of Round5 INDCCA public-key encryption, for NIST security levels 1, 3 and 5. No forward

error correction is used, hence the 0 at the end of the parameter designator. All polynomial multiplications are done modulo the prime-order cyclotomic polynomial. This choice is considered more conservative since it uses the same design principles as Round2.

R5N1_{1,3,5}KEM_0c: Parameters for the non-ring/unstructured variant ($n = 1$) of Round5 key-encapsulation mechanism, for NIST security levels 1, 3 and 5. No forward error correction is used, hence the 0 at the end of the parameter designator. Since the generation of the public-parameter $\boldsymbol{A}$ is expensive, we include performance results for three alternative ways of generating it denoted as $T0, T1, T2$ and discuss performance trade-offs between the usage of SHAKE256 and AES128.

R5N1_{1,3,5}PKE_0c: Parameters for the non-ring/unstructured variant ($n = 1$) of Round5 public-key encrpytion, for NIST security levels 1, 3 and 5. No forward error correction is used, hence the 0 at the end of the parameter designator.

1.9.2 Round5 parameters for specific use-cases

Round5 further details three additional "specific use-case" parameters with the only purpose of demonstrating the flexibility of Round5 and applicability to three different specialized usage scenarios.

R5ND_0KEM_2iot: A small parameter set targeting the Internet of Things use-cases, with lower bandwidth and computational footprints, yet still providing 90 bits of (quantum) security. XE2 forward error correction is used to improve failure rate, bandwidth and security.

R5ND_1KEM_4longkey: An alternative to theNIST level 1, ring variant key-encapsulation parameter set R5ND_1KEM that encapsulates a 192-bit key despite targeting the NIST security level 1, to ensure that the (quantum) cost of atacking the encapsulated key (e.g., by using Grover's quantum search algorithm) is as much as the (quantum) cost of attacking the underlying cryptosystem, i.e., Round5.

R5N1_3PKE_0smallCT: An alternative to the NIST level 3, non-ring variant public-key encryption parameter set R5N1_3PKE_0smallCT that has exceptionally small ciphertexts, targeting usage scenarios where the public-key is static and hence exchanged rarely, implying that bandwidth footprint depends on the size of the ciphertext. Hence this parameter set, despite enjoying the more conservative security assumption based on unstructured lattices, has a bandwidth requirement comparable to ring or structured variants.

1.9.3 Implementations

We include three implementations of Round5.

- Reference: This is a reference implementation of Round5 capable to run all configurations at run time. Matrix and polynomial operations are explicit so that it is simple to verify the correctness of the implementation.

- Optimized: This is an optimized implementation of Round5 capable to run all configurations at run time. This implementation shows the feasibility of an implementation running very different configurations with a common code base.

- Tiny: This is an optimized implementation of Round5 capable of running a single configuration, fixed at compiler time. This implementation can also fix – at compile time – the way to generate $\boldsymbol{A}$, i.e., $f_{d,n}^{(0)}$, $f_{d,n}^{(1)}$, or $f_{d,n}^{(2)}$ and whether to use SHAKE or AES. Trade-offs are discussed in the following sections.

1.9.4 Development Environment

The performance numbers or Round5 have been gathered on a MacBookPro10.1 with an Intel Core i7 2.6GHz, running macOS 10.14.1. The code has been compiled with `gcc -march=native -mtune=native -O3 -fomit-frame-pointer -fwrapv`, using Apple LLVM version 10.0.0 (clang-1000.11.45.5).

All tests were run 10000 times, the measurements shown are the minimum values of all test runs and are for the tiny implementation of the algorithm. We use the minumum values as opposed to say the median or average since we believe the minimum is a better illustration of the time actually spent by the algorithm and not induced by e.g., system or multi-tasking overhead.

For the memory requirements, we have provided an indication based on a possible implementation. Note that the actual memory usage depends, of course, on the specifics of a particular implementation (e.g., an implementation might require matrices to exist (also) in transposed form, need room for storing intermediate results, etc.).

1.9.5 Round5 CPA_KEM: Parameters

This section contains specific parameter values for Round5.KEM fitting the Round5 configurations described at the beginning of this section.

Tables 2, 3, and 4 show the configurations for the Round5 parameter sets for the NIST security levels 1, 3, and 5. Table 5 present the configurations for the parameter sets for specific use cases.

In all of the above tables, the values under *Security Levels* represent the cost of known attacks (discussed in Section 1.7) against the underlying LWR or RLWR problem.

The main conclusions from these tables are as follows:

- All parameter sets strictly fulfill NIST security levels and the hybrid attack is the one that has the highest impact on Round5.

- Parameter sets using XE5 (R5ND_1,3,5KEM_5c) have around 25% lower bandwidth requirements compared with those without error correction (R5ND_1,3,5KEM_5c). Since Round5 proposes a KEM that offers CPA security, no active attacks are applicable.

- The application specific IoT parameter set (R5ND_0KEM_2iot) only requires 736 bytes and it still offers a classical security level of 129 bits if an enumeration cost model is applied.

- Many security protocols that require a handshake to exchange a session key do not require active security and a relatively high failure probability is good enough. This allows finding parameters that offer smaller keys and ciphertext at the price of a slightly higher failure rate. For instance, for Internet of Things deployment the failure rate of a wireless communication link will be likely worse than the failure rate of R5ND_0KEM_2iot, and thus, a failure probability of 2^{-41} is sufficient. Alternatively, it is possible to further increase this failure probability so that the sizes of public-key and ciphertext further drop.

Table 2: R5ND_{1,3,5}KEM_0c parameters

	Security Level		
	NIST1	**NIST3**	**NIST5**
Parameters			
$d/n/h$	618/618/104	786/786/384	1018/1018/428
$q/p/t/b$	$2^{11}/2^8/2^4/2^1$	$2^{13}/2^9/2^4/2^1$	$2^{14}/2^9/2^4/2^1$
$\overline{n}/\overline{m}$	1/1	1/1	1/1
$f/xe/\kappa/\mu$	0/0/128/128	0/0/192/192	0/0/256/256
Performance			
Total bandwidth (bytes)	1316	1890	2452
Public key (bytes)	634	909	1178
Ciphertext (bytes)	682	981	1274
Failure rate	2^{-65}	2^{-71}	2^{-64}
Classical/Quantum security (bits) − Core Sieving			
Primal attack	144/131	192/175	256/233
Dual attack	147/134	193/175	258/235
Hybrid attack	128/122	194/183	261/248
Sparse-secrets attack	144/130	192/174	256/232
Classical/Quantum optimal block size BKZ − Core Sieving			
Primal attack	494/494	659/659	878/878
Dual attack	504/504	662/662	885/885
Hybrid attack	437/461	665/691	895/934
Classical/Quantum security (bits) − Enumeration			
Primal attack	340/170	500/250	729/365
Dual attack	350/175	503/252	737/368
Hybrid attack	160/133	293/222	397/309
Sparse-secrets attack	340/170	500/250	729/364
Classical/Quantum optimal block size BKZ − Enumeration			
Primal attack	494/494	659/659	878/878
Dual attack	504/504	661/662	885/885
Hybrid attack	284/412	442/603	554/773

Table 3: R5ND_{1,3,5}KEM_5c parameters

	Security Level		
	NIST1	**NIST3**	**NIST5**
Parameters			
$d/n/h$	490/490/162	756/756/242	940/940/414
$q/p/t/b$	$2^{10}/2^7/2^3/2^1$	$2^{12}/2^8/2^2/2^1$	$2^{12}/2^8/2^2/2^1$
$\overline{n}/\overline{m}$	1/1	1/1	1/1
$f/xe/\kappa/\mu$	5/190/128/318	5/218/192/410	5/234/256/490
Performance			
Total bandwidth (bytes)	994	1639	2035
Public key (bytes)	445	780	972
Ciphertext (bytes)	549	859	1063
Failure rate	2^{-88}	2^{-117}	2^{-64}
Classical/Quantum security (bits) − Core Sieving			
Primal attack	130/118	194/176	256/233
Dual attack	132/120	197/179	259/235
Hybrid attack	128/122	193/183	263/249
Sparse-secrets attack	130/118	194/176	256/232
Classical/Quantum optimal block size BKZ − Core Sieving			
Primal attack	446/446	666/666	878/878
Dual attack	452/452	674/674	887/887
Hybrid attack	440/462	660/691	900/940
Classical/Quantum security (bits) − Enumeration			
Primal attack	297/148	507/254	729/365
Dual attack	301/151	515/258	739/369
Hybrid attack	170/135	270/215	390/307
Sparse-secrets attack	296/148	507/253	729/364
Classical/Quantum optimal block size BKZ − Enumeration			
Primal attack	446/446	666/666	878/878
Dual attack	451/451	673/674	887/887
Hybrid attack	297/416	416/588	547/770

Table 4: R5N1_{1,3,5}KEM_0c parameters

	Security Level		
	NIST1	**NIST3**	**NIST5**
Parameters			
$d/n/h$	594/1/238	881/1/238	1186/1/712
$q/p/t/b$	$2^{13}/2^{10}/2^7/2^3$	$2^{13}/2^{10}/2^7/2^3$	$2^{15}/2^{12}/2^7/2^4$
$\overline{n}/\overline{m}$	7/7	8/8	8/8
$f/xe/\kappa/\mu$	0/0/128/43	0/0/192/64	0/0/256/64
Performance			
Total bandwidth (bytes)	10450	17700	28552
Public key (bytes)	5214	8834	14264
Ciphertext (bytes)	5236	8866	14288
Failure rate	2^{-66}	2^{-65}	2^{-77}
Classical/Quantum security (bits) – Core Sieving			
Primal attack	132/121	201/184	257/234
Dual attack	131/120	202/184	256/233
Hybrid attack	128/121	192/182	256/241
Sparse-secrets attack	130/119	201/183	256/233
Classical/Quantum optimal block size BKZ – Core Sieving			
Primal attack	422/422	658/658	847/847
Dual attack	418/418	659/659	843/844
Hybrid attack	408/423	626/652	844/871
Classical/Quantum security (bits) – Enumeration			
Primal attack	275/138	499/250	696/348
Dual attack	272/136	500/250	692/346
Hybrid attack	177/130	268/210	420/304
Sparse-secrets attack	271/135	499/249	691/345
Classical/Quantum optimal block size BKZ – Enumeration			
Primal attack	422/422	658/658	847/847
Dual attack	418/418	659/659	843/843
Hybrid attack	306/404	414/577	578/765

Table 5: Specific use case Round5 KEM parameters

	Parameter Set	
	R5ND_0KEM_2iot	**R5ND_1KEM_4longkey**
Parameters		
$d/n/h$	372/372/178	490/490/162
$q/p/t/b$	$2^{11}/2^7/2^3/2^1$	$2^{10}/2^7/2^3/2^1$
$\overline{n}/\overline{m}$	1/1	1/1
$f/xe/\kappa/\mu$	2/53/128/181	4/163/192/355
Performance		
Total bandwidth (bytes)	736	1016
Public key (bytes)	342	453
Ciphertext (bytes)	394	563
Failure rate	2^{-41}	2^{-71}
Classical/Quantum security (bits) – Core Sieving		
Primal attack	98/89	130/118
Dual attack	98/89	132/120
Hybrid attack	96/90	128/122
Sparse-secrets attack	97/88	130/118
Classical/Quantum optimal block size BKZ – Core Sieving		
Primal attack	335/335	446/446
Dual attack	336/336	452/452
Hybrid attack	329/341	440/462
Classical/Quantum security (bits) – Enumeration		
Primal attack	201/100	297/148
Dual attack	202/101	301/151
Hybrid attack	129/96	170/135
Sparse-secrets attack	200/100	296/148
Classical/Quantum optimal block size BKZ – Enumeration		
Primal attack	335/335	446/446
Dual attack	336/336	451/451
Hybrid attack	243/325	297/416

1.9.6   Round5 CPA_KEM: CPU and Memory Requirements

This section contains the CPU and memory requirements for Round5 CPA_KEM described in Section 1.9.5.

Table 6 shows the performance and memory usage figures for the tiny implementation of the R5ND_{1,3,5}KEM_0c algorithm. Table 7 shows the performance and memory usage figures for the tiny implementation of the R5ND_{1,3,5}KEM_5c algorithm. Table 8 shows the performance and memory usage figures for the tiny implementation of the R5N1_{1,3,5}KEM_0c $\tau = 2$ algorithm. Table 9 shows the performance of the parameter sets for specific use cases of the tiny implementation of the Round5 CPA_KEM algorithm. Table 10 compares the performance of the different variants of computing $\boldsymbol{A}$ of the tiny implementation of the R5N1_1KEM_0c algorithm.

The main conclusions from these tables are as follows:

- CPU performance for ring configurations does not seem to be an issue for most applications since it requires fractions of a millisecond without using any type of hardware instructions.

- CPU performance for non-ring configurations can be sufficient for many applications, except those that require very high throughput.

- The performance of R5N1_1KEM_0c for different choices of $f_{d,n}^{(\tau)}$ implemented using as DRBG either SHAKE128 or AES128 with hardware acceleration is illustrated. We observe that CPU performance of $f_{d,n}^{(0)}$ is worse than $f_{d,n}^{(2)}$ that is also worse than $f_{d,n}^{(1)}$. The reason is simple: $f_{d,n}^{(1)}$ requires the smallest amount of pseudorandom data to be obtained from the DRBG while $f_{d,n}^{(0)}$ requires the highest amount of pseudorandom data. $f_{d,n}^{(2)}$ is in between and has the lowest memory cost since the amount of data to be kept in memory is limited. Finally, we observe that if the CPU has hardware-enabled AES instructions, then AES-based generation of $\boldsymbol{A}$ is competitive.

Table 6: Tiny R5ND_{1,3,5}KEM_0c performance and memory usage

	Security Level		
	NIST1	**NIST3**	**NIST5**
API Parameters			
CRYPTO_SECRETKEYBYTES	16	24	32
CRYPTO_PUBLICKEYBYTES	634	909	1,178
CRYPTO_BYTES	16	24	32
CRYPTO_CIPHERTEXTBYTES	682	981	1,274
Performance: Elapsed time (ms)			
KEM Generate Key Pair	0.01	0.03	0.05
KEM Encapsulate	0.01	0.04	0.06
KEM Decapsulate	0.007	0.02	0.02
Total	0.03	0.10	0.1
Performance: CPU Clock Cycles			
KEM Generate Key Pair	29.1K	88.1K	118.1K
KEM Encapsulate	37.4K	114.7K	149.3K
KEM Decapsulate	17.3K	49.6K	57.8K
Total	83.8K	252.3K	325.1K
Memory usage indication			
KEM Generate Key Pair	4,374B	5,673B	7,350B
KEM Encapsulate	7,784B	10,182B	12.9KiB
KEM Decapsulate	3,458B	4,581B	5,954B

Table 7: Tiny R5ND_{1,3,5}KEM_5c performance and memory usage

	Security Level		
	NIST1	**NIST3**	**NIST5**
API Parameters			
CRYPTO_SECRETKEYBYTES	16	24	32
CRYPTO_PUBLICKEYBYTES	445	780	972
CRYPTO_BYTES	16	24	32
CRYPTO_CIPHERTEXTBYTES	549	859	1,063
Performance: Elapsed time (ms)			
KEM Generate Key Pair	0.01	0.02	0.04
KEM Encapsulate	0.02	0.03	0.05
KEM Decapsulate	0.01	0.01	0.03
Total	0.05	0.07	0.1
Performance: CPU Clock Cycles			
KEM Generate Key Pair	32.8K	56.7K	100.7K
KEM Encapsulate	56.0K	80.7K	136.8K
KEM Decapsulate	32.5K	38.6K	69.5K
Total	121.3K	175.9K	307.0K
Memory usage indication			
KEM Generate Key Pair	3,417B	5,364B	6,676B
KEM Encapsulate	6,182B	9,631B	11.7KiB
KEM Decapsulate	2,813B	4,339B	5,431B

Table 8: Tiny R5N1_{1,3,5}KEM_0c $\tau = 2$ performance and memory usage

	Security Level		
	NIST1	**NIST3**	**NIST5**
API Parameters			
CRYPTO_SECRETKEYBYTES	16	24	32
CRYPTO_PUBLICKEYBYTES	5,214	8,834	14,264
CRYPTO_BYTES	16	24	32
CRYPTO_CIPHERTEXTBYTES	5,236	8,866	14,288
Performance: Elapsed time (ms)			
KEM Generate Key Pair	0.6	1.0	3.8
KEM Encapsulate	0.8	1.3	5.2
KEM Decapsulate	0.07	0.09	0.3
Total	1.5	2.4	9.2
Performance: CPU Clock Cycles			
KEM Generate Key Pair	1,631K	2,643K	9,826K
KEM Encapsulate	2,097K	3,432K	13.4M
KEM Decapsulate	171.8K	228.7K	686K
Total	3,900K	6,303K	23.9M
Memory usage indication			
KEM Generate Key Pair	38.5KiB	53.9KiB	117.4KiB
KEM Encapsulate	51.9KiB	76.6KiB	150.1KiB
KEM Decapsulate	21.5KiB	36.4KiB	51.2KiB

Table 9: Tiny Round5 CPA_KEM specific use case parameters, performance and memory usage

	Parameter Set	
	R5ND_0KEM_2iot	**R5ND_1KEM_4longkey**
API Parameters		
CRYPTO_SECRETKEYBYTES	16	24
CRYPTO_PUBLICKEYBYTES	342	453
CRYPTO_BYTES	16	24
CRYPTO_CIPHERTEXTBYTES	394	563
Performance: Elapsed time (ms)		
KEM Generate Key Pair	0.01	0.01
KEM Encapsulate	0.02	0.02
KEM Decapsulate	0.01	0.02
Total	0.04	0.06
Performance: CPU Clock Cycles		
KEM Generate Key Pair	31.3K	35.0K
KEM Encapsulate	49.9K	59.1K
KEM Decapsulate	36.1K	48.1K
Total	117.3K	142.1K
Memory usage indication		
KEM Generate Key Pair	2,606B	3,441B
KEM Encapsulate	4,744B	6,348B
KEM Decapsulate	2,186B	2,979B

Table 10: Tiny R5N1_1KEM_0c – $\boldsymbol{A}$ generation variants

	Variants			
	$\tau = 0$	$\tau = 0$	$\tau = 1$	$\tau = 2$
	(cSHAKE128)	(AES128-HW)	(cSHAKE128)	(cSHAKE128)
Performance: Elapsed time (ms)				
KEM Generate Key Pair	2.2	0.6	0.5	0.6
KEM Encapsulate	2.5	0.8	0.8	0.8
KEM Decapsulate	0.07	0.07	0.07	0.07
Total	4.8	1.5	1.4	1.5
Performance: CPU Clock Cycles				
KEM Generate Key Pair	5,829K	1,590K	1,412K	1,631K
KEM Encapsulate	6,364K	2,129K	2,021K	2,097K
KEM Decapsulate	171.7K	176.5K	171.7K	171.8K
Total	12.4M	3,896K	3,605K	3,900K
Memory usage indication				
KEM Generate Key Pair	710.5KiB		2,088.8KiB	38.5KiB
KEM Encapsulate	723.9KiB		2,102.2KiB	51.9KiB
KEM Decapsulate	21.5KiB		21.5KiB	21.5KiB

1.9.7   Round5 CCA_PKE: Parameters

This section contains specific parameter values for Round5.PKE fitting the Round5 configurations described at the beginning of this section.

Tables 2, 3, and 4 show the configurations for the Round5 parameter sets for the NIST security levels 1, 3, and 5. Table 5 present the configurations for the parameter sets for specific use cases.

In all of the above tables, the values under *Security Levels* represent the cost of known attacks (discussed in Section 1.7) against the underlying LWR or RLWR problem.

The main conclusions from these tables are as follows:

- As for r5_cpa_kem parameters, all parameter sets strictly fulfil NIST security levels and the hybrid attack is the one that has the highest impact on Round5.

- Parameter sets using XE5 (R5ND_1,3,5PKE_5c) have around 25% lower bandwidth requirements compared with those parameters without error correction (R5ND_1,3,5KEM_5c). For instance, R5ND_5PKE_5c requires 730 Bytes less than R5ND_5PKE_0c.

- r5_cca_pke, targeting IND-CCA security, requires lower failure probability compared with r5_cpa_kem; this leads to parameters that require slightly longer public-keys and ciphertexts. When no error correction is used, this difference is larger and can be up to 400 Bytes.

- The application specific R5N1_3PKE_0smallCT, i.e., the parameter set for a non-ring configuration with short ciphertext is a desirable alternative in use cases in which the public-key remains static for a long period of time since the ciphertext is small and comparable in size with ring configurations, minimizing the overall bandwidth requirement.

Table 11: R5ND_{1,3,5}PKE_0c parameters

	Security Level		
	NIST1	**NIST3**	**NIST5**
Parameters			
$d/n/h$	586/586/182	852/852/212	1170/1170/222
$q/p/t/b$	$2^{13}/2^9/2^4/2^1$	$2^{12}/2^9/2^5/2^1$	$2^{13}/2^9/2^5/2^1$
$\overline{n}/\overline{m}$	1/1	1/1	1/1
$f/xe/\kappa/\mu$	0/0/128/128	0/0/192/192	0/0/256/256
Performance			
Total bandwidth (bytes)	1432	2102	2874
Public key (bytes)	676	983	1349
Encryption overhead (bytes)	756	1119	1525
Failure rate	2^{-155}	2^{-147}	2^{-143}
Classical/Quantum security (bits) − Core Sieving			
Primal attack	131/119	199/181	281/255
Dual attack	132/120	202/183	286/260
Hybrid attack	128/121	192/182	257/246
Sparse-secrets attack	130/118	199/180	281/255
Classical/Quantum optimal block size BKZ − Core Sieving			
Primal attack	448/448	683/683	963/963
Dual attack	451/451	692/692	981/981
Hybrid attack	439/457	657/688	880/929
Classical/Quantum security (bits) − Enumeration			
Primal attack	298/149	524/262	822/411
Dual attack	301/151	534/267	842/421
Hybrid attack	177/135	266/213	346/289
Sparse-secrets attack	298/149	524/262	822/411
Classical/Quantum optimal block size BKZ − Enumeration			
Primal attack	448/448	683/683	963/963
Dual attack	451/451	692/692	981/981
Hybrid attack	306/417	412/584	500/736

Table 12: R5ND_{1,3,5}PKE_5c parameters

	Security Level		
	NIST1	**NIST3**	**NIST5**
Parameters			
$d/n/h$	508/508/136	756/756/242	940/940/414
$q/p/t/b$	$2^{10}/2^7/2^4/2^1$	$2^{12}/2^8/2^3/2^1$	$2^{12}/2^8/2^3/2^1$
$\overline{n}/\overline{m}$	1/1	1/1	1/1
$f/xe/\kappa/\mu$	5/190/128/318	5/218/192/410	5/234/256/490
Performance			
Total bandwidth (bytes)	1097	1730	2144
Public key (bytes)	461	780	972
Encryption overhead (bytes)	636	950	1172
Failure rate	2^{-142}	2^{-256}	2^{-144}
Classical/Quantum security (bits) – Core Sieving			
Primal attack	133/121	194/176	256/233
Dual attack	135/122	197/179	259/235
Hybrid attack	128/122	193/183	263/249
Sparse-secrets attack	132/120	194/176	256/232
Classical/Quantum optimal block size BKZ – Core Sieving			
Primal attack	455/455	666/666	878/878
Dual attack	462/462	674/674	887/887
Hybrid attack	437/462	660/691	900/940
Classical/Quantum security (bits) – Enumeration			
Primal attack	305/152	507/254	729/365
Dual attack	311/156	515/258	739/369
Hybrid attack	166/134	270/215	390/307
Sparse-secrets attack	304/152	507/253	729/364
Classical/Quantum optimal block size BKZ – Enumeration			
Primal attack	455/455	666/666	878/878
Dual attack	462/462	673/674	887/887
Hybrid attack	292/413	416/588	547/770

Table 13: R5N1_{1,3,5}PKE_0c parameters

	Security Level		
	NIST1	**NIST3**	**NIST5**
Parameters			
$d/n/h$	636/1/114	876/1/446	1217/1/462
$q/p/t/b$	$2^{12}/2^9/2^6/2^2$	$2^{15}/2^{11}/2^7/2^3$	$2^{15}/2^{12}/2^9/2^4$
$\overline{n}/\overline{m}$	8/8	8/8	8/8
$f/xe/\kappa/\mu$	0/0/128/64	0/0/192/64	0/0/256/64
Performance			
Total bandwidth (bytes)	11544	19392	29360
Public key (bytes)	5740	9660	14636
Encryption overhead (bytes)	5804	9732	14724
Failure rate	2^{-146}	2^{-144}	2^{-144}
Classical/Quantum security (bits) − Core Sieving			
Primal attack	146/133	194/177	257/234
Dual attack	146/134	193/176	257/234
Hybrid attack	128/122	192/181	256/241
Sparse-secrets attack	145/133	192/175	257/234
Classical/Quantum optimal block size BKZ − Core Sieving			
Primal attack	469/469	632/632	848/848
Dual attack	471/471	628/628	847/847
Hybrid attack	407/428	626/646	844/874
Classical/Quantum security (bits) − Enumeration			
Primal attack	317/159	473/237	697/348
Dual attack	319/160	469/235	696/348
Hybrid attack	159/130	296/214	402/300
Sparse-secrets attack	317/158	469/234	695/347
Classical/Quantum optimal block size BKZ − Enumeration			
Primal attack	469/469	632/632	848/848
Dual attack	470/471	628/628	847/847
Hybrid attack	283/404	445/586	559/756

Table 14: Specific use case Round5 PKE parameters

	Parameter Set R5N1_3PKE_0smallCT
Parameters	
$d/n/h$	757/1/378
$q/p/t/b$	$2^{14}/2^9/2^4/2^1$
$\overline{n}/\overline{m}$	192/1
$f/xe/\kappa/\mu$	0/0/192/192
Performance	
Total bandwidth (bytes)	164524
Public key (bytes)	163536
Encryption overhead (bytes)	988
Failure rate	2^{-149}
Classical/Quantum security (bits) − Core Sieving	
Primal attack	194/177
Dual attack	193/176
Hybrid attack	191/181
Sparse-secrets attack	192/175
Classical/Quantum optimal block size BKZ − Core Sieving	
Primal attack	632/632
Dual attack	628/628
Hybrid attack	624/648
Classical/Quantum security (bits) − Enumeration	
Primal attack	473/237
Dual attack	469/235
Hybrid attack	281/211
Sparse-secrets attack	469/234
Classical/Quantum optimal block size BKZ − Enumeration	
Primal attack	632/632
Dual attack	628/628
Hybrid attack	428/580

1.9.8   Round5 CCA_PKE: CPU and Memory Requirements

This section contains the CPU and memory requirements for Round5 CCA_PKE described in Section 1.9.7.

Table 15 shows the performance and memory usage figures for the tiny implementation of the R5ND_{1,3,5}PKE_0c algorithm. Table 16 shows the performance and memory usage figures for the tiny implementation of the R5ND_{1,3,5}PKE_5c algorithm. Table 17 shows the performance and memory usage figures for the tiny implementation of the R5N1_{1,3,5}PKE_0c $\tau = 2$ algorithm. Table 18 shows the performance of the parameter sets for specific use cases of the tiny implementation of the Round5 CCA_PKE algorithm. Some conclusions from these tables are as follows:

- Due to the asymmetry of $\bar{n}$ and $\bar{m}$, the key generation time of R5N1_3PKE-_0smallCT is relatively high; however, this is only a one-time cost since public-keys remain static for a long period of time. On the other hand, the benefit of the asymmetry is that the encryption time decreases.

- Ring versions of Round5 perform currently around 70 times faster than its non-ring versions.

Table 15: Tiny R5ND_{1,3,5}PKE_0c performance and memory usage

	Security Level		
	NIST1	NIST3	NIST5
API Parameters			
CRYPTO_SECRETKEYBYTES	708	1,031	1,413
CRYPTO_PUBLICKEYBYTES	676	983	1,349
CRYPTO_BYTES	756	1,119	1,525
Performance: Elapsed time (ms)			
PKE Generate Key Pair	0.02	0.02	0.03
PKE Encrypt	0.03	0.04	0.05
PKE Decrypt	0.03	0.05	0.06
Total	0.07	0.1	0.1
Performance: CPU Clock Cycles			
PKE Generate Key Pair	41.9K	59.8K	80.7K
PKE Encrypt	63.6K	95.5K	124.6K
PKE Decrypt	81.5K	118.5K	151.8K
Total	187.1K	273.8K	357.1K
Memory usage indication			
PKE Generate Key Pair	4,916B	7,150B	9,814B
PKE Encrypt	7,596B	10.8KiB	14.8KiB
PKE Decrypt	4,112B	6,014B	8,226B

Table 16: Tiny R5ND_{1,3,5}PKE_5c performance and memory usage

	Security Level		
	NIST1	**NIST3**	**NIST5**
API Parameters			
CRYPTO_SECRETKEYBYTES	493	828	1,036
CRYPTO_PUBLICKEYBYTES	461	780	972
CRYPTO_BYTES	636	950	1,172
Performance: Elapsed time (ms)			
PKE Generate Key Pair	0.01	0.02	0.04
PKE Encrypt	0.02	0.04	0.06
PKE Decrypt	0.03	0.05	0.08
Total	0.07	0.1	0.2
Performance: CPU Clock Cycles			
PKE Generate Key Pair	32.9K	57.3K	101.1K
PKE Encrypt	60.7K	95.1K	156.1K
PKE Decrypt	83.3K	124.5K	215.8K
Total	176.8K	276.9K	473.0K
Memory usage indication			
PKE Generate Key Pair	4,018B	6,168B	7,680B
PKE Encrypt	6,481B	9,746B	11.9KiB
PKE Decrypt	3,465B	5,258B	6,576B

Table 17: Tiny R5N1_{1,3,5}PKE_0c $\tau = 2$ performance and memory usage

	Security Level		
	NIST1	**NIST3**	**NIST5**
API Parameters			
CRYPTO_SECRETKEYBYTES	5,772	9,708	14,700
CRYPTO_PUBLICKEYBYTES	5,740	9,660	14,636
CRYPTO_BYTES	5,804	9,732	14,724
Performance: Elapsed time (ms)			
PKE Generate Key Pair	0.4	2.0	2.7
PKE Encrypt	0.5	2.6	3.6
PKE Decrypt	0.5	2.7	3.7
Total	1.4	7.3	10.0
Performance: CPU Clock Cycles			
PKE Generate Key Pair	930K	5,139K	7,036K
PKE Encrypt	1,307K	6,684K	9,230K
PKE Decrypt	1,385K	7,005K	9,541K
Total	3,622K	18.8M	25.8M
Memory usage indication			
PKE Generate Key Pair	40.4KiB	112.0KiB	133.1KiB
PKE Encrypt	50.6KiB	126.0KiB	152.4KiB
PKE Decrypt	31.4KiB	46.6KiB	67.0KiB

Table 18: Tiny Round5 CCA_PKE specific use case parameters, performance and memory usage

	Parameter Set R5N1_3PKE_0smallCT $\tau = 2$
API Parameters	
CRYPTO_SECRETKEYBYTES	163,584
CRYPTO_PUBLICKEYBYTES	163,536
CRYPTO_BYTES	988
Performance: Elapsed time (ms)	
PKE Generate Key Pair	30.9
PKE Encrypt	1.2
PKE Decrypt	3.2
Total	35.2
Performance: CPU Clock Cycles	
PKE Generate Key Pair	80.0M
PKE Encrypt	2,993K
PKE Decrypt	8,192K
Total	91.2M
Memory usage indication	
PKE Generate Key Pair	920.7KiB
PKE Encrypt	481.8KiB
PKE Decrypt	446.5KiB

1.10   Advantages and limitations

Flexibility in the underlying problem's choice Round5 can be configured to rely on the Learning with Rounding (LWR) or Ring-Learning with Rounding (RLWR) problems, by controlling the input parameters n and d. This allows the user to flexibly choose the configuration (and underlying problem instantiation) that fits best his application and security requirements, even while Round5 is already in deployment. For instance, a user dealing with strictly confidential information might only trust a public-key encryption (PKE) algorithm with a construction having no additional (e.g., ring) structure, while another user operating a wireless network for a less critical application would prioritize low bandwidth and/or energy requirements and thus might prefer a (more efficient) ring-based construction. The unified approach provides from day one a contingency and smooth transition path, should vulnerabilities in ring-based constructions be discovered, since the non-ring based construction is already deployed. Finally, as the Round5 implementation and code remains unified independent of the underlying problem instantiated, the amount of code in devices and effort required for code-review is reduced.

Small public-keys and ciphertexts Round5 relies on rounding (specifically, the GLWR problem, see Section 1.3), leading to public-keys and ciphertexts with coefficients that have only $\lceil \log p \rceil$ and $\lceil \log t \rceil$ bits. Furthermore, Round5 relies on prime cyclotomic polynomials that provide a large pool of (q, n) values to choose from, allowing the selection of parameters that satisfy security requirements while minimizing bandwidth requirements. The usage of contant-time XEf error correction allows dealing with multiple errors so that even smaller parameters $(n = d, p)$, and thus, messages are feasible. Round5 thus is suited for bandwidth-constrained applications, and Internet protocols that only allow limited packet sizes. Comparing with other submissions, Round5 shows a very good trade-off between security and the sizes of public keys and ciphertexts, as reported in [50] (with "Performance" as label for the X-axis).

Common building blocks for KEM and PKE By design, r5_cpa_kem and r5_cca_pke are constructed using common building blocks. This allows for a common security and correctness analysis for both schemes. Furthermore, it reduces the amount of code in devices, and the effort required for code-review of r5_cpa_kem and r5_cca_pke.

Flexibility in achieving security levels The design choices in Round5, especially the choice for prime cyclotomic polynomials, allowed the fine-tuning of parameters to each NIST level. Round5 thus enables the user to choose parameters that tightly fit the required security level.

Flexibility for bandwidth Different applications can have different security needs and operational constraints. Round5 enables the flexible choice of parameters so that it is possible to adjust bandwidth requirements and

security level according to the application needs. Round5 achieves this by using prime cyclotomic polynomials and rounding. For instance, configuration R5ND_1KEM_5c for NIST security level 1 requires the exchange of 994 Bytes so that it can be used by an application with communication constraints; configuration R5N1_5PKE_0c offers a much higher security level and does not rely on any ring structure so that it might be a preferred option for the exchange of highly confidential information. The amount of data to be exchanged for the latter configuration (29360 Bytes) is larger than for the former, but is smaller than in another existing non-ring proposal for an equivalent security level [25].

Flexibility for CPU Different applications can have different security needs and operational constraints. Round5 allows for flexibility in the choice of parameters so that it is possible to adjust CPU needs and security requirements according to the application needs. Furthermore, the parameters in Round5 are chosen such that a unified implementation performs well for any input value (n, d). If necessary, optimized implementations for specific parameters can be realized, leading to even faster operation than with the unified implementation. For instance, configuration R5N1_5KEM_5c for NIST security level 5, intended for a high security level and not relying on a ring structure, requires around 9.2 milliseconds for key generation, encapsulation and decapsulation in our testing platform. Another application, requiring faster operation and with less security needs, can use configuration R5ND_1KEM_5c for NIST security level 1 that performs a factor 70 faster.

Flexibility for cryptographic primitives Round5 and its building blocks can be used to create cryptographic primitives such as authenticated key-exchange schemes in a modular way, e.g., as in [26].

Flexibility for integration into real-world security protocols The Internet relies on security protocols such as TLS, IKE, SSH, IPsec, DNSSEC etc. Round5 can be easily integrated into them because it has relatively small messages and is computationally efficient. For instance, the public-key and ciphertext in R5ND_5KEM_5c for NIST security level 5 require a total of 2035 Bytes. This is smaller than other lattice-based proposals such as [26] or [7], making sure that it fits in the protocols without requiring changes or additional complexity such as packet fragmentation. At the same time, all Round5 configurations can be realized by means of a single library minimizing the amount of work to extend and maintain them.

Prevention of pre-computation and back-door attacks Round5 offers different alternatives to refresh $\boldsymbol{A}$ in a way that is efficient but also secure against pre-computation and back-door attacks. In the ring setting, this is achieved by computing a new $\boldsymbol{A}$. In the non-ring setting, three options are provided: $(f_{d,n}^{(0)})$ randomly generating $\boldsymbol{A}$ in each protocol exchange,

$(f_{d,n}^{(1)})$ permuting a fixed and system-wide public parameter $\boldsymbol{A_{master}}$ in each protocol exchange, or $(f_{d,n}^{(2)})$ deriving $\boldsymbol{A}$ from a large pool of values. – determined in each protocol exchange – by means of a permutation. Permutation-based approaches show a performance advantage compared with generating $\boldsymbol{A}$ randomly— for instance, in settings in which a server has to process many connections. This is because the server can keep a fixed $\boldsymbol{A}$ in memory and just permute it according to the client request. If keeping a fixed $\boldsymbol{A}$ in memory in the client is an issue, then $f_{d,n}^{(2)}$ has a clear advantage compared with $f_{d,n}^{(1)}$ due to its lower memory needs.

Post-deployment flexibility A single implementation can be used for all Round5 configuration parameters, without the need to recompile the code. This approach reduces the amount of code in the devices and makes code review easier. Furthermore, it enables a unified Round5 deployment that can be customized to fit non-ring or ring-based use cases, and a smooth transition to non ring usage, should vulnerabilities in ring-based constructions be discovered.

Efficiency in constrained platforms The implementation of Round5 uses up to 16 bit long integers to represent the Round5 data. This enables good performance even in architectures with constrained processors. In particular, all parameter sets using error correction have parameter $p \leq 2^8$ so that this parameter set is specially suitable for resource constrained platforms. Round5 incorporates the special parameter set R5ND_0KEM_2iot that has a total bandwidth requirement of just 736 Bytes, while providing a reasonable security level (129-bits classical enumeration). This is feasible due to the usage of prime cyclotomic rings that allows configuring Round5 with a $n = d$ parameter between 256 and 512, the only two options available with power-of-two cyclotomic rings.

Parallelization Operations in Round5, for instance matrix multiplications, can be parallelized allowing for faster performance. This type of optimization has not been applied yet.

Resistance against side channel attacks The design of Round5 allows for efficient constant-time implementations, since by design all secrets are ternary and have a fixed number of ones and minus ones. The XEf error correction codes avoid conditions altogether and table look-ups and are therefore resistant against timing attacks.

Low failure probability The failure probability for all proposals for r5_cpa_kem (except the one for the IoT use case) is at most 2^{-64}. The failure probability for all proposals for r5_cca_pke is at most 2^{-142}. These failure probabilities can be achieved because of the usage of sparse, ternary secrets and the large pool of parameter sets. The usage of XEf error correction codes in some Round5 parameter sets plays a fundamental role to deal with multiple errors and decrease the overall failure probability.

Known underlying mathematical problems Round5 builds on the well-known Learning with Rounding and Ring Learning with Rounding problems.

Provable security We provide security reductions from the sparse ternary LWR and RLWR problems to r5_cpa_kem and r5_cca_pke, and from the standard Learning with Errors (LWE) problem to the sparse ternary LWR problem. Even though the latter reduction is not tight, this gives confidence in the Round5 design. As the parameters using error correction require performing certain operations modulo $x^{n+1} - 1$, the above security reductions do not apply to them. We describe an RLWE-based analogon of Round5 with error correction, and present a reduction for such a scheme to RLWE. This give confidence in Round5 with error correction. We discuss why this reduction does not directly translate to the RLWR case, and argue that it does not have any impact on the concrete security estimates.

Easy adoption A device that supports Round5 can handle a wide range of configuration parameters, for both the ring and the non-ring versions, without re-compilation. This flexibility will ease a smooth adoption.

Perfect forward secrecy Round5's key generation algorithm is fast, which makes it suitable for scenarios in which it is necessary to refresh a public/private key pair in order to maintain forward secrecy.

Application-specific configurations The flexibility of Round5 allows addressing the needs of a very different range of applications. Round5 includes a ring-based KEM configuration addressing *Internet of Things* that achieves very small bandwidth requirements at the price of lower security, still providing $129-$bit security (classical enumeration). Another special configuration is a ring-based KEM NIST level 1 configuration in which the encapsulated key is 192-bit long instead of just 128-bit, so that the difficulty of attacking the encapsulated key (by Grover) is approximately equal to the difficulty of quantum lattice attack to Round5. This configuration can be very practical for many applications and has a total bandwidth requirement of just 1016 Bytes. Finally, Round5 includes a non-ring-based PKE parameter set with small encryption overhead (988 Bytes) at the cost of a larger public key. The latter configuration makes unstructured lattice configurations more efficient in applications in which the public-key remains static, e.g., email security.

1.11   Technical Specification of Reference Implementation

This section specifies Round5 from an implementation perspective. The description is close to Round5's C reference implementation and aims at explaining it and allowing other developers to create own implementations. In the description, we avoid special symbols to facilitate its direct implementation. For in-

stance, Greek letters are not included and, e.g., we write $kappa$ instead of κ; we write q_bits instead of q_{bits} or n_bar instead of $\bar{n}$.

1.11.1   Round5 main parameters

Each Round5 parameter set (Sections 1.9.1 and 1.9.2) is determined by the following parameters that are located on the top of each parameter table in Sections 1.9.5 and 1.9.7.

d	Dimension of underlying lattice.
n	If ring-based, $n > 1$. If non-ring-based, $n = 1$.
h	Number of non-zero values per column in secret matrices.
q, q_bits	Power of two modulo size where $q = 2^{q\_bits}$.
p, p_bits	Power of two modulo size where $p = 2^{p\_bits}$.
t, t_bits	Power of two modulo size where $t = 2^{t\_bits}$.
b, b_bits	b_bits bits are extracted per ciphertext symbol; $b = 2^{b\_bits}$.
n_bar	Number of columns of the secret matrix to compute $\boldsymbol{B}$.
m_bar	Number of columns of the secret matrix to compute $\boldsymbol{U}$.
$kappa$	Security parameter; number of information bits in error-correcting code.
f	Number of bit errors correctable by error-correcting code.
xe	Number of parity bits of error correcting code.
mu	Number of ciphertext symbols: $mu \cdot b\_bits \geq kappa + xe$.

1.11.2   Round5 derived or configurable parameters

The following parameters are derived from the main parameters. Parameter tau indicates a configuration choice in function $f_{d,n}^{(\tau)}$ used to generate $\boldsymbol{A}$.

d/n	Number of polynomial elements in a row of $\boldsymbol{A}$.
$kappa\_bytes$	Security parameter in bytes, i.e., kappa/8.
z	$z = \max(p, \frac{tq}{p})$. Relevant in reduction proofs and definition of rounding constants.
z_bits	$z\_bits = \lceil log2(z) \rceil$.
h_1	$h\_1 = \frac{q}{2p}$ is a rounding constant.
h_2	$h\_2 = \frac{q}{2z}$ is a rounding constant.

h_3	$h\_3 = \frac{p}{2t} + \frac{p}{2b} - \frac{q}{2z}$ is a constant to reduce decryption bias.
h_4	$h\_4 = \frac{q}{2p} - \frac{q}{2z}$ is a constant to reduce decryption bias.
N	Identifies polynomial $N_{n+1}(x) = x^{n+1} - 1$.
Phi	Identifies polynomial $\Phi_{n+1}(x) = N_{n+1}(x)/(x-1)$.
Xi	Identifies if reduction polynomial is N or Phi.
tau	Index in $f_{d,n}^{(\tau)}$.

For all proposed Round5 parameter sets, $z = p$, $h\_1 = h\_2 = \frac{q}{2p}$, and $h\_4 = 0$.

1.11.3 Basic data types, functions and conversions

Bit	A binary digit: 0 or 1.
Bit string	An ordered sequence of bits, e.g., $[0, 1, 1, 1]$ contains 4 bits, 0 is the left-most bit. Bit strings are denoted by a lower case letter.
Byte	An ordered sequence of 8 bits, e.g., $[0, 1, 1, 1, 0, 1, 1, 1]$. Bytes are interpreted in little-endian. The previous example represents $0xEE$ in hexadecimal or 238 in decimal representation.
Byte string	An ordered sequence of bytes e.g., $[0x01, 0x02, 0x03]$ contains 3 bytes, 0x01 is the left-most byte. Byte strings are denoted by a lower case letter.
0^c	A bit string containing c 0s.
$\|\|$	The concatenation operator is represented by $\|\|$ and concatenates two strings, e.g., two bit or two byte stings. For instance, $$[0, 1, 1, 1, 0, 1] = [0, 1, 1]\|\|[1, 0, 1]$$
$\lceil a \rceil$	For a real number a, the ceiling operator is represented by $\lceil a \rceil$ and returns the next integer value of a. For instance, $\lceil 15/8 \rceil = 2$, $\lceil -3.2 \rceil = -3$.
$\lfloor a \rfloor$	For a real number a, the flooring operator is represented by $\lfloor a \rfloor$ and returns the integer part of value a. For instance, $\lfloor 15/8 \rfloor = 1$, $\lfloor -3.2 \rfloor = -4$.
$\oplus$	For two bits x and y, $x \oplus y$ represents the modulo 2 addition of two bits. For instance, $1 \oplus 1 = 0$.

Little endianness

Round5 uses a little endianness representation for bit strings and byte strings as defined above as well as for vectors and polynomials. In a bit string this means that the leftmost bit is the least significant bit and the rightmost bit or the last bit is the most significant bit. The same applies to byte string, vectors, and polynomials.

For instance, the following byte c contains 8 bits, starting with bit 0 (c_0) and ending with bit 7 (c_7).

$$c = [c_0, c_1, c_2, c_3, c_4, c_5, c_6, c_7]$$

In another example, a 64-bit word w contains 8 bytes such as the first byte is byte 0 (c_0) and the last byte is byte 7 (c_7).

$$w = [c_0, c_1, c_2, c_3, c_4, c_5, c_6, c_7]$$

For instance, $c = [0, 1, 0, 1, 1, 1, 0, 1]$ equals 186, or $0xba$ in hexadecimal representation.
$w = [0xef, 0xcd, 0xab, 0x89, 0x67, 0x45, 0x23, 0x01]$ equals $0x0123456789abcdef$.

Moduli

Round5 performs operations modulo r with $r = \{q, p, t\}$ being powers of two. The number of bits is denoted as $r\_bits = \{q\_bits, p\_bits, t\_bits\}$ such that $r = 2^{r\_bits}$.

Vectors

Round5 vectors $\boldsymbol{v}_r^{(k,1)}$ are defined column-wise, and contain k integers in $\mathbb{Z}_r$. Vector element with index 0 $\boldsymbol{v}[0]$ comes first (little-endian) followed by element with index 1 $\boldsymbol{v}[1]$ till last element with index $k - 1$, i.e., $\boldsymbol{v}[k - 1]$. A vector element $\boldsymbol{v}[i]$ is represented as a bit string containing $r\_bits = \lceil log_2 r \rceil$ bits. Vectors are denoted by a bold low case letter.

Polynomials

Round5 polynomials are in $\mathcal{R}_{n,r} = \mathbb{Z}_r[x]/(\xi_{n+1}(x))$ where $r = \{q, p, t\}$ is a power of two, $n+1$ is chosen to be a prime number, and $\xi_{n+1}(x)$ can be either the prime cyclotomic polynomial $\Phi(x) = 1 + x + \ldots + x^n$, or the NTRU polynomial $N_{n+1}(x) = x^{n+1} - 1 = (x-1) \cdot \Phi_{n+1}(x)$ having $n+1$ coefficients. All coefficients are $r\_bits = \lceil log_2(r) \rceil$ bits long, i.e., $r\_bits = \{q\_bits, p\_bits, t\_bits\}$. When $n = 1$, polynomials have a single coefficient in $\mathbb{Z}_r$.

Polynomials are represented in little endian, i.e., coefficient of degree 0 comes first and the highest degree coefficient comes last.

A polynomial is represented as a vector $\boldsymbol{poly}$ so that $\boldsymbol{poly}[i]$ is the coefficient of degree i. In particular,

$$\boldsymbol{poly} = \boldsymbol{poly}[0] + \boldsymbol{poly}[1] \cdot x + \cdots + \boldsymbol{poly}[s] \cdot x^s$$

where $s = n$ when $Xi = N$ and $s = n-1$ when $Xi = Phi$.

Each polynomial coefficient $\boldsymbol{poly}[i]$ is represented as a bit string of length r_bits.

$$\boldsymbol{poly}[i] = [\boldsymbol{poly}[i]\_0, \boldsymbol{poly}[i]\_1, \ldots, \boldsymbol{poly}[i]\_(r\_bits - 1)]$$

Polynomial vectors	Round5 polynomial vectors $\boldsymbol{v}_r^{(k,1)}$ are defined column-wise, and contain $k = d/n$ polynomials in $\mathcal{R}_{n,r} = \mathbb{Z}_r[x]/\left(\xi_{n+1}(x)\right)$. Polynomial vectors are denoted by a bold lower case letter.
Matrices	Round5 matrices $\boldsymbol{A}_r^{(row,col)}$ contain row rows and col columns and their elements are polynomials in $\mathcal{R}_{n,r} = \mathbb{Z}_r[x]/\left(\xi_{n+1}(x)\right)$. Matrices are denoted by a bold capital letter.

1.11.4   Supporting functions

SHAKE	Specified in FIPS-202, SHAKE with security levels 128 and 256 is used as the extendable-output function to generate the coefficients in public parameter $\boldsymbol{A}$ and secrets $\boldsymbol{S}$ and $\boldsymbol{R}$.
cSHAKE	Specified in NIST SP 800-185, cSHAKE128 and cSHAKE256 offer security levels 128 and 256 and build on top of SHAKE128 and SHAKE256 such that $cSHAKE(X, L, N = "", S = "") = SHAKE(X, L)$ where X, L, N and S are as in NIST SP 800-185. $cSHAKE$ allows for customization bit strings so that $cSHAKE(X1, L1, N1, S1)$ and $cSHAKE(X1, L1, N2, S2)$ produce unrelated outputs unless $N1 = N2$ and $S1 = S2$.
	$cSHAKE$ is used in the construction of the permutations specified in $f_{d,n}^{(1)}$ and $f_{d,n}^{(2)}$.
	Round5 operations are specified in terms of $cSHAKE$ although some operations can be replaced by $SHAKE$ as $cSHAKE(X, L, N = "", S = "") = SHAKE(X, L)$.
AES	Specified in FIPS-197, AES with security levels 128, 192 and 256 is used as symmetric block cipher in Round5. AES is also used as an optional alternative to generate random data.
randombytes_init	Function to initialize a $true$ random number generator $randombytes()$.

randombytes()   Function that returns B bytes of true random data de-
                                 noted as *randomness*.

$$randomness = randombytes(B)$$

hash()   Function to obtain an *output_len* byte long hash *output* of
                                 an *input_len* byte long *input* using a *customization_string*
                                 of a given length.
                                 *output = hash(output_len, input, input_len, customiza-
                                 tion_string, customization_string_len)*
                                 Round5 uses cSHAKE with a security strength of 128-
                                 or 256-bits as the default hash function due to its fast
                                 performance in software and the requirement to obtain
                                 hashes of variable length. If the S customization_string
                                 equals "", then Round5 hash is equivalent to SHAKE
                                 with a security strength of 128- or 256-bits:

$$output = hash(output\_leninput, input\_len)$$

                                 Usage for different key lengths is specified in Table 19.

drbg_init_customization()
                                 Function to initialize a Deterministic Random Bit Gen-
                                 erator (DRBG) that takes as input a seed *seed* and a
                                 customization_string. This function also *hides* complex-
                                 ity of $f_{d,n}^{(\tau)}$ and its implementation using either cSHAKE
                                 or AES.

$$drbg\_init(seed, customization\_string)$$

                                 Round5 uses cSHAKE as the default DRBG due to its
                                 fast performance in software. When cSHAKE is used,
                                 drbg_init is implemented with cSHAKE_Absorb with a
                                 security strength of 128- or 256-bits.

$$drbg\_init = cSHAKE\_absorb(seed, "", customization\_string)$$

                                 Round5 also supports AES as an optional DRBG due
                                 to its fast performance in hardware. When AES is used,
                                 R5_DRBG_Init involves two steps. First, it derives $seed_{AES}$
                                 using *cSHAKE()*

$$seed\_AES = hash(seed\_len, seed, seed\_len, customization\_string, 2)$$

                                 then it uses *seed_AES* to initialize the AES key with
                                 security level of *seed_len* bits. Usage for different key
                                 lengths is specified in Table 19.

drbg_init()   Function used instead of *drbg_init_customization()* when *customization_string* equals "". This function has an identical definition, but it already assumes *customization_string* = "", and thus, the usage of cSHAKE is equivalent to the direct usage of SHAKE.

drbg()   Function to obtain a *len*-byte long byte string denoted as *randomdata* from Round5's Deterministic Random Bit Generator (DRBG).

$$randomdata = drbg(len)$$

Round5 uses cSHAKE as the default DRBG due to its fast performance in software. When cSHAKE is used, Round5's DRBG implements cSHAKE_Squeeze with a security strength of 128- or 256-bits.

$$drbg(len) = cSHAKE\_Squeeze(len)$$

Round5 also supports AES as an optional DRBG due to its fast performance in hardware. When AES is used, Round5's DRBG applies AES in counter mode (NIST SP 800-38D) for a security level of 128-, 192- or 256-bits with input 16 byte long block of all zeros $0x00000000000000000000000000000000$.

$$drbg(len) = AESCTR(len)$$

Usage for different key lengths is specified in Table 19.

The retrieved bits from the underlying random bit function must always be interpreted as little endian. Thus, they will need to be reversed on a big endian machine.

drbg_sampler16   Function to sample random numbers uniformly distributed in a given range.

$$randomdata = drbg\_range16(range)$$

This function requires initializing the drbg with *drbg_init* before its first invocation. It first computes the greatest positive integer *range_divisor* such that $range\_divisor \cdot range < 2^{16}$. Then it uses $drbg(2^{16})$ to compute random numbers x till $x < range\_divisor \cdot range$. The returned value is $\lfloor x/range\_divisor \rfloor$

drbg_sampler16_2   Function to sample random numbers uniformly distributed in a given range that is a power of two.

$$randomdata = drbg\_range16\_2(range)$$

This function requires initializing the drbg with $drbg\_init$ before its first invocation. It uses $drbg(2^{16})$ to compute a 2 byte random number x. The returned value corresponds to the first $r\_bits = log2(range)$ bits of x, i.e., $[x\_0, x\_1, \ldots, x\_(r\_bits - 1)]$ For instance, if $drbg$ returns [0x12, 0x34] and $range = 2^{12}$, then the sampled element is [0x12, 0x30].

1.11.5   Cryptographic algorithm choices

Round5 includes a drbg, a function to generate pseudorandom data, a hash function, and a DEM. These functions are implemented by means of cSHAKE (or optionally, AES in counter mode), and AES in GCM mode. The main reason for choosing cSHAKE is the use the customization string in the generation of the permutation in $f_{d,n}^{(\tau)}$. In other cases, the customization string is not included and in those cases cSHAKE is equivalent to SHAKE. The reason for including – optionally – AES in counter mode is its fast performance.

Table 19: Cryptographic algorithm choices

	$kappa$	drbg_init	drbg	hash	DEM
	128	cSHAKE128_absorb	cSHAKE128_squeeze	cSHAKE128	GCM-AES128
Def.	192	cSHAKE256_absorb	cSHAKE256_squeeze	cSHAKE256	GCM-AES192
	256	cSHAKE256_absorb	cSHAKE256_squeeze	cSHAKE256	GCM-AES256
	128	hash	CTR-AES-128	cSHAKE128	GCM-AES128
		Init CTR-AES128			
	192	hash	CTR-AES-192	cSHAKE256	GCM-AES192
Opt.		Init CTR-AES192			
	256	hash	CTR-AES-256	cSHAKE256	GCM-AES256
		Init CTR-AES256			

The security strength of the algorithms is determined as a function of the length of the secret, $kappa$. Since cSHAKE is also defined with security strengths of 128− and 256− bits, cSHAKE128 is used when $kappa = 128$; otherwise, cSHAKE256 is used.

The default configuration (Def.) means that those are the default choices in Round5 code. This choice is motivated by the faster performance of cSHAKE in software compared with AES. The optional configuration (Opt.) means that those are optional choices in Round5 code. This choice is motivated by the fact that some platforms have hardware-enabled AES instructions that increase the performance of the pseudo random data generation.

1.11.6 Core functions

mult_poly_ntru   Function that multiplies input polynomials $\boldsymbol{a}$ and $\boldsymbol{b}$ in $\mathcal{R}_{n,r} = \mathbb{Z}_r[x]/\left(N_{n+1}(x)\right)$ obtaining polynomial $\boldsymbol{c}$.

$$\boldsymbol{c} = mult\_poly\_ntru(\boldsymbol{a}, \boldsymbol{b}, n, r)$$

lift_poly   Function that lifts input polynomial $\boldsymbol{a}$ in $\mathbb{Z}_r[x]/\left(\Phi_{n+1}(x)\right)$ obtaining polynomial $\boldsymbol{a^{(L)}}$ in $\mathbb{Z}_r[x]/\left(N_{n+1}(x)\right)$.

$$\boldsymbol{a^{(L)}} = lift\_poly(\boldsymbol{a}, n, r)$$

In mathematical terms, lifting means:

$$\boldsymbol{a^{(L)}} = (x-1)\boldsymbol{a}$$

Coefficient-wise, this means that:

$$\boldsymbol{a}^{(L)} = \boldsymbol{a}_0^{(L)} + \boldsymbol{a}_1^{(L)} \cdot x + \boldsymbol{a}_2^{(L)} \cdot x^2 + \cdots + \boldsymbol{a}_n^{(L)} x^n$$

equals

$$-\boldsymbol{a}_0 + (\boldsymbol{a}_0 - \boldsymbol{a}_1)x + (\boldsymbol{a}_1 - \boldsymbol{a}_2)x^2 + \cdots + (\boldsymbol{a}_{n-2} - \boldsymbol{a}_{n-1})x^{n-1} + \boldsymbol{a}_{n-1}x^n$$

unlift_poly   Function that unlifts input polynomial $\boldsymbol{a^{(L)}}$ in $\mathbb{Z}_r[x]/\left(N_{n+1}(x)\right)$, obtaining $\boldsymbol{a}$ in $\mathbb{Z}_r[x]/\left(\Phi_{n+1}(x)\right)$.

$$\boldsymbol{a} = unlift\_poly(\boldsymbol{a^{(L)}}, n, r)$$

In mathematical terms, unlifting means:

$$\boldsymbol{a} = \boldsymbol{a^{(L)}}/(x-1)$$

As $\boldsymbol{a^{(L)}} = lift_poly(\boldsymbol{a}, n, r)$, the coefficients of $\boldsymbol{a}$ can be recursively computed as

$$\boldsymbol{a}_0 = -\boldsymbol{a}_0^L \text{ and } \boldsymbol{a}_i = \boldsymbol{a}_{i-1} - \boldsymbol{a}_i^L \text{ for } 1 \leq i \leq n-1.$$

mult_poly   Function that takes input polynomials $\boldsymbol{a}$ and $\boldsymbol{b}$ of length n and multiplies them in $\mathcal{R}_{n,r} = \mathbb{Z}_r[x]/\left(\xi_{n+1}(x)\right)$ obtaining polynomial $\boldsymbol{c}$.

$$\boldsymbol{c} = mult\_poly(\boldsymbol{a}, \boldsymbol{b}, n, r, Xi)$$

If $Xi = N$, then $\boldsymbol{c} = mult\_poly\_ntru(\boldsymbol{a}||0, \boldsymbol{b}, n, r)$. If $Xi = Phi$, then this function lifts input parameter $\boldsymbol{a}$ before calling $mult\_poly\_ntru$ and unlifts the result $\boldsymbol{c}$, i.e., $\boldsymbol{c} = unlift(mult\_poly\_ntru(lift(\boldsymbol{a}, n, r), \boldsymbol{b}, n, r), n, r)$.

add_poly   Function that takes input polynomials $\boldsymbol{a}$ and $\boldsymbol{b}$ of length n and adds them component-wise obtaining $\boldsymbol{c}$.

$$\boldsymbol{c} = add\_poly(\boldsymbol{a}, \boldsymbol{b}, n, r, Xi)$$

mult_matrix   Function that multiplies two matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ of dimension $A\_row \times A\_col$ and $B\_row \times B\_col$, with elements in $\mathcal{R}_{n,r} = \mathbb{Z}_r[x]/(\xi_{n+1}(x))$ obtaining matrix $\boldsymbol{C}$ of dimension $A\_row \times B\_col$ whose elements are also in $\mathcal{R}_{n,r}$.

$$\boldsymbol{C} = mult\_matrix(\boldsymbol{A}, A\_row, A\_col, \boldsymbol{B}, B\_row, B\_col, n, r, Xi)$$

```
1  for(i = 0; i < A_row; i + +)
2      for(j = 0; j < B_col; j + +)
3          C[i, j] = 0
4          for(k = 0; k < A_col; k + +)
5              tmp = mult_poly(A[i, k], B[k, j], n, r, Xi)
6              C[i, j] = add_poly(C[i, j], tmp, n, r)
```

transpose_matrix   Function that transposes input matrix $\boldsymbol{A}$ of dimension $A\_row \times A\_col$ with elements in $\mathcal{R}_{n,r} = \mathbb{Z}_r[x]/(\Phi_{n+1}(x))$ obtaining matrix $\boldsymbol{A\_T}$ of dimension $A\_col \times A\_row$ whose elements are also in $\mathcal{R}_{n,r}$.

$$\boldsymbol{A\_T} = transpose\_matrix(\boldsymbol{A}, A\_row, A\_col, n, r)$$

round_element   Function that rounds a_bits bits long element x from a_bits to b_bits bits using rounding constant h where $a = 2^{a\_bits}$ and $b = 2^{b\_bits}$.

$$x = round(x, a\_bits, b\_bits, h) = \left\lfloor \frac{x + h}{2^{a\_bits - b\_bits}} \right\rfloor = \left\lfloor \frac{b}{a}(x+h) \right\rfloor$$

round_matrix   Function that performs coefficient-wise rounding applying *round_element* to all polynomial coefficients in matrix $\boldsymbol{A}$ of dimension $A\_row \times A\_col$ with elements in $\mathcal{R}_{n,2^{a\_bits}} = \mathbb{Z}_{2^{a\_bits}}[x]/(\xi_{n+1}(x))$, obtaining matrix $\boldsymbol{B}$ of dimension $A\_row \times A\_col$ with elements in $\mathcal{R}_{n,2^{b\_bits}} = \mathbb{Z}_{2^{b\_bits}}[x]/(\xi_{n+1}(x))$.

$$\boldsymbol{B} = round\_matrix(\boldsymbol{A}, A\_row \cdot A\_col, n, a\_bits, b\_bits, h)$$

decompress   Function that decompresses element x from b_bits to a_bits bits .

$$x = decompress(x, b\_bits, a\_bits) = x \cdot 2^{a\_bits - b\_bits}$$

decompress_ matrix   Function that performs coefficient-wise decompression using *decompress* of the first *mu* polynomial coefficients in matrix $\boldsymbol{A}$ of dimension $A\_row \times A\_col$ with elements in $\mathcal{R}_{n,2^{b\_bits}} = \mathbb{Z}_{2^{b\_bits}}[x]/\left(\xi_{n+1}(x)\right)$, obtaining vector $\boldsymbol{b}$ of dimension *mu* with elements in $\mathbb{Z}_{2^{a\_bits}}$.

$$\boldsymbol{b} = decompress\_matrix(\boldsymbol{A}, mu, 1, b\_bits, a\_bits)$$

permutation_tau_1   Function that computes permutation $\boldsymbol{p1}$ associated to $f_{d,1}^{(1)}$ used to permute the row elements in a fixed master public parameter matrix $\boldsymbol{A\_master}$ and obtain the public parameter A.

$$\boldsymbol{p1} = permutation\_tau\_1(sigma)$$

$\boldsymbol{p1}$ is a vector of length d with elements in $\mathbb{Z}_d$. permutation_tau_1($sigma$) is obtained by first initializing the DRBG with *drbg_init(seed, customization_string=0x0001)* and then sampling the elements with *drbg_sampler16()* with range d.

permutation_tau_2   Function that computes permutation $\boldsymbol{p2}$ associated to $f_{d,1}^{(2)}$ used to permute the elements in a master public parameter vector $\boldsymbol{a\_master}$ and obtain the public parameter A.

$$\boldsymbol{p2} = permutation\_tau\_2(sigma)$$

$\boldsymbol{p2}$ is a vector of length d with elements in $\mathbb{Z}_q$. Permutation_tau_2(sigma) is obtained by first initializing the DRBG with *drbg_init(seed, customization_string=0x0001)* and then running *drbg_sampler16_2* with range q.

create_A   Function that computes master public parameter $\boldsymbol{A}$ given random *sigma*.

$$\boldsymbol{A} = create\_A(sigma)$$

Internally, depending on the choice of *tau*, create_A computes $\boldsymbol{A}$ using one out of three strategies that provide a trade-off between CPU and memory performance (see Section 1.4.2).

$\mathbf{f_{d,n}^{(0)}}$: Applies *drbg_sampler16_2* with range q. This function must be initialized with seed *sigma* and without customization string. Output of *drbg_sampler16_2* is used to fill-in $\boldsymbol{A}$ row-wise, i.e., first polynomial element 0 in row

0, then polynomial element 1 in row 0, till all polynomial elements in row 0 are assigned; then polynomial element 0 in row 1, then polynomial element 1 in row 1, till all polynomial elements in row 1 are assigned. This process goes on till all elements in $\boldsymbol{A}$ are initialized. Each element in $\boldsymbol{A}$ is a polynomial, and it is filled-in coefficient-wise, i.e., first coefficient of degree 0, then coefficient of degree 1, till the coefficient of highest degree.

When $n = 1$, $\boldsymbol{A}$ consists of d^2 polynomial elements in $\mathcal{R}_{n,q} = \mathbb{Z}_q$. Since all Round5 parameter sets are such that $2^{16} \geq q > 2^8$, a total of d^2 calls to *drbg_sampler16_2* are required. When $n = d$, $\boldsymbol{A}$ consists of a single element in $\mathcal{R}_{n,r} = \mathbb{Z}_q[x]/\left(\Phi_{n+1}(x)\right)$. Since all Round5 parameter sets are such that $2^{16} \geq q > 2^8$, thus, a total of d calls to *drbg_sampler16_2* are required. In both cases, ($n = 1$ and $n = d$) element i is assigned the output of the i^{th} call to *drbg_sampler16_2*.

$\mathbf{f_{d,n}^{(1)}}$: only applies to $n = 1$. This requires that a matrix $\boldsymbol{A\_master}$ of size $d \times d$ has been precomputed and is a public-parameter known to all parties in the system. It requires permutation vector $\boldsymbol{p1}$ – computed with *permutation_tau_1()* – of length d with elements randomly chosen in $\mathbb{Z}_d$ to permute the elements in each row in A_master to obtain $\boldsymbol{A}$, in particular: $\boldsymbol{A}[i,j] = \boldsymbol{A\_master}[i, (j + \boldsymbol{p1}[j]) \bmod \text{d}]$.

The reference and optimized implementations fill in $\boldsymbol{A\_master}$ from a precomputed matrix $\boldsymbol{A\_fixed}$ of the same dimensions. This is done to enable an implementation that can run ring and non-ring parameters with a different *tau* choice.

$\mathbf{f_{d,n}^{(2)}}$: only applies to $n = 1$. This approach first computes a vector $\boldsymbol{a}_{master}$ of length q by calling q times *drbg_sampler16_2* with range q. This function must be initialized with seed *sigma* and without customization string.

It then uses a permutation vector $\boldsymbol{p2}$ computed with *permutation_tau_2* of length d with elements $\mathbb{Z}_q$ to pickup d consecutive elements in a_master and construct $\boldsymbol{A}$ in this way, in particular:
$\boldsymbol{A}[i,j] = \boldsymbol{a}_{master}[(\boldsymbol{p2}[i] + j) \bmod \text{q}]$

create_ secret_vector   Function that computes a sparse ternary secret $\boldsymbol{s}$ of di-

mension $length \times 1$ and having fixed hamming weight h.

$$s = create\_secret\_vector(length, h)$$

Each secret vector s contains d/n polynomials in $\mathcal{R}_{n,r} = \mathbb{Z}_r[x]/(\xi_{n+1}(x))$. Since n can only take values 1 and d, then this secret vector can represent two different data structures always consisting of d elements in $\mathbb{Z}_r$:

- ring case: a single polynomial containing n ternary coefficients, with exactly $h/2$ "+1" and $h/2$ "−1" values.

- non-ring case: d polynomials, each of them having a single ternary coefficient, and the vector with exactly $h/2$ "+1" and $h/2$ "−1" values.

Since in both data structures, there are exactly $h/2$ "+1" and $h/2$ "−1" values, then this function only computes the h positions in vector s containing the h non-zero elements. To this end, this function uses $drbg\_sample16$ to sample the h non-zero positions with $i = 0, ..., h − 1$. This is done by checking whether a position $s[i]$ is occupied already, or not. The i^{th} sampled position is assigned value "−1" if i is odd and "+1" if i is even as described below.

```
1  s[length] = 0
2  for(i = 0; i < h; i + +)
3    do
4        x = drbg_sampler16(d)
5    while(s[x]! = 0);
6    if(is_even(i))
7      s[x] = 1;
8    else
9      s[x] = −1;
```

create_S_T   Function that computes secret S^T of dimension $\bar{n} \times d$ and hamming weight h per row given function $create\_secret\_vector(d, h)$.

$$\boldsymbol{S\_T} = create\_S\_T(\bar{n}, length = d/n{\cdot}n, hamming\_weight = h)$$

create_R_T   Function that computes secret R^T of dimension $\bar{m} \times d$ and hamming weight h per row given function $create\_secret\_vector(d, h)$.

$$\boldsymbol{R\_T} = create\_R\_T(\bar{m}, length = d/n{\cdot}n, hamming\_weight = h)$$

sample_mu   Function that returns the first mu polynomial coefficients of $\boldsymbol{M}$ where $\boldsymbol{M}$ contains $\bar{n} \times \bar{m}$ polynomials in $\mathcal{R}_{n,r} = \mathbb{Z}_r[x]/(\xi_{n+1}(x))$.

When $n = d$, $\boldsymbol{M}$ contains a single polynomial $\boldsymbol{m}$ and the values returned by *sample_mu* corresponds to the first mu polynomial coefficients:

$$m0 + m1x + \cdots + m_{mu-1}x^{mu-1}$$

When $n = 1$, $\boldsymbol{M}$ contains $\bar{n} \times \bar{m}$ polynomials, each of them having a single element in $\mathbb{Z}_r$ denoted as $M_{i,j}$ where i and j refer to the row and column indexes. The returned coefficients are the first mu coefficients in the matrix iterated row-wise.

$$\overbrace{\underbrace{M_{0,0}, M_{0,1}, \ldots, M_{0,\bar{n}-1}}_{\text{Row } 0}, \ldots \underbrace{M_{ii-1,0}, M_{ii-1,1}, \ldots, M_{ii-1,jj-1}}_{\text{Row } ii-1}}^{mu \text{ coefficients}},$$

where $mu = ii \cdot \bar{n} + jj$

add_msg   Function that converts a binary string message m of length $mu \cdot b\_bits$ into mu elements in $\mathbb{Z}_t$ and then adds them to a vector $\boldsymbol{x}$ of length mu whose elements are each t_bits long.

$$\boldsymbol{result} = add\_msg(len, \boldsymbol{x}, m, b\_bits, t\_bits)$$

Addition of message bits $m[i \cdot b\_bits, \ldots, (i+1) \cdot b\_bits - 1]$ is done by first multiplying this bit string interpreted as an element in $\mathbb{Z}_b$ by t/b so that this is represented in $\mathbb{Z}_t$.

For instance, if $t = 2^5$ and $b = 2^2$, and $m[0, 1] = [1, 1]$, then its bit representation in $\mathbb{Z}_t$ is $[0,0,0,1,1]$ where the left most bit is bit 0 and the right most bit is bit 4. This value is added to the t_bits-bit value $\boldsymbol{x}[0]$.

diff_msg   Function that computes difference of mu elements in *vecv* and $\boldsymbol{x}$ modulo p.

$$\boldsymbol{result} = diff\_msg(len, \boldsymbol{v}, \boldsymbol{x}, p)$$

xef_compute   Round5's error correction operates in a bit string m of size $mu \cdot b\_bits$, of which $kappa$ bits are used to transport a shared secret message of length 128-, 192- or 256- bits and the remaining $xe = mu - kappa$ bits are used to correct errors in it. The XEf design is easily scalable, so

that it is possible to use the parameter f to correct a variable number of *errors* in a message.

$$m = xef\_compute(m, \kappa, f)$$

is the function that given the κ bits long message computes the xe parity bits consisting of the values of $2f$ registers. The lengths $\{l_0, l_1, \ldots, l_{2f-1}\}$ of registers $\{r_0, r_1, \ldots, r_{2f-1}\}$ for $XE(\kappa, f)$ are as follows:

- $XE(128, 2) : \{11, 13, 14, 15\}$
- $XE(192, 4) : \{13, 15, 16, 17, 19, 23, 29, 31\}$
- $XE(128, 5) : \{16^{(*)}, 11, 13, 16, 17, 19, 21, 23, 25, 29\}$
- $XE(192, 5) : \{24^{(*)}, 13, 16, 17, 19, 21, 23, 25, 29, 31\}$
- $XE(256, 5) : \{16^{(*)}, 16, 17, 19, 21, 23, 25, 29, 31, 37\}$

The lengths of $xe = \sum l_i$ for $XE(128, 2)$, $XE(192, 4)$, $XE(128, 5)$, $XE(192, 5)$, $XE(256, 5)$ are 53, 163, 190, 218, and 234, respectively.

Bit j in register r_i of length l_i is computed as:

$$r_i[j] = m[j] \oplus m[j + l_i] \oplus \cdots \oplus m[j + \lfloor \frac{\kappa - 1 - j}{l_i} \rfloor \cdot l_i]$$

As in HILA5, register r_0 in XE5 codes is special – marked with $(*)$ – and computes the register bits as the block-wise XOR of kappa/l_0 consecutive string bits.

$$r_0[j] = m[j \cdot \frac{\kappa}{l_0}] \oplus \cdots \oplus m[(j + 1) \cdot \frac{\kappa}{l_0} - 1]$$

Given a code $XE(\kappa, f)$, a κ bit message m and the set r including $2f$ registers $r_0, r_1, \ldots, r_{2f-1}$ obtained from m and with total length xe, the output bit string of $xef\_compute$ is:

$$m \oplus [0^{kappa} || r_0 || r_1 || \ldots || r_{2f-1}]$$

xef_correct

Function that corrects up to f errors distributed over the block of $mu \cdot b\_bits = kappa + xe$ bits.

$$m = xef\_correct(m, \kappa, f)$$

This function requires applying $xe\_compute$ on the received bit string of length $mu \cdot b\_bits$ before error correction. When doing so, the registers r_i'' computed from the

received message m' are XORed with the received registers r'_i , i.e., $r'''_i = r'_i \oplus r''_i$ so that only the bit differences are left in the last xe bits of the received bit string.

$$\underbrace{m'[0,\ldots,kappa-1]}_{kappa \text{ bits}} || \underbrace{r'''_0||\ldots||r'''_{2f-1}}_{xe \text{ bits}} = compute\_xef(m',kappa,f)$$

Message bit $\boldsymbol{m'}[k]$ is flipped if:

$$\sum_{i=0}^{2f-1} r'''_i[f_{r_i}(k)] > f$$

where $f_{r_i}(j)$ determines the bit in register r_i that depends on message bit j as specified in xef_compute depending on whether r_i is a special register or not.

pack	Function that serializes an input vector $\boldsymbol{input}$ containing $input\_length$ components, each of size $input\_size$ bits and places it into an output message $output$ of size $\lceil input\_length \cdot input\_size/8 \rceil$ bytes.

$$output = pack(\boldsymbol{input}, input\_length, input\_size)$$

For instance, input vector $[0x8, 0x1, 0xf, 0x2]$ contains 4 elements, each of them 4 bits long is serialized in output vector $[0x81, 0xf2]$. If the size is not a multiple of 8 bits, then the most significant bits of the last byte are padded with 0s.

pack_pk Function that serializes the components $sigma$ and $\boldsymbol{B}$ of Round5 public key. The function is defined as follows and uses $pack$ internally.

$$pk = pack\_pk(sigma, ss\_size, \boldsymbol{B}, d/n \cdot n\_bar \cdot n, p\_bits)$$

By definition, the $sigma$ is packed in the first ss_size bytes of pk. Next, the $d/n \times \bar{n}$ n-coefficient polynomials in $\boldsymbol{B}$ are packed. Packing is done row-wise (first element 0 in row 0, then element 1 in row 0,..., till the last element in row 0; next the second row till the last row). For each polynomial element, packing is done starting with coefficient of degree 0 and ending with the coefficient of degree $n-1$. For each polynomial coefficient having p_bits bits, bit 0 comes first and bit p_bits-1 comes last.

pack_ct Function that serializes the components $\boldsymbol{U}$ and $\boldsymbol{v}$ of Round5 ciphertext. The function is defined as follows and uses $pack$ internally.

$$ct = pack\_ct(\boldsymbol{U}, d/n \cdot m\_bar \cdot n, p\_bits, \boldsymbol{v}, mu, t\_bits)$$

By definition, $\boldsymbol{U}$ is packed in the first $\lceil d/n \cdot m\_bar \cdot n \cdot p\_bits/8 \rceil$ bytes of ct. Packing is done row-wise (first element 0 in row 0, then element 1 in row 0,..., till the last element in row 0; next the second row till the last row). For each polynomial element, packing is done starting with coefficient of degree 0 and ending with the coefficient of degree $n - 1$. For each polynomial coefficient having p_bits bits, bit 0 comes first and bit $p\_bits - 1$ comes last. Next, v is packed packing first $v[0]$, then $v[1]$, till $v[mu - 1]$. Each coefficient $v[i]$ of v has t_bits bits and those are also packed following a little endian approach, i.e., first bit 0, then bit 1, till bit $t\_bits - 1$.

unpack

Function that deserializes a bit string *input* of length $\lceil number\_elements \cdot element\_bits/8 \rceil$ bytes into a vector of size $number\_elements$ in which each vector element is in $\mathbb{Z}_{2^{element\_bits}}$.

$\boldsymbol{output} = unpack(input, number\_elements, element\_bits)$

Bits $\{i * element\_bits, \dots, (i + 1) * element\_bits - 1\}$ in the input bit string correspond to element $\boldsymbol{output}[i]$ of the deserialized output vector.

unpack_pk

Function that de-serializes the components *sigma* and $\boldsymbol{B}$ of Round5's public-key *pk*. The function is defined as follows and uses *unpack* internally.

$(sigma, \boldsymbol{B}) = unpack\_pk(pk, sigma\_size, B\_elements, B\_element\_bits)$

where the input parameters are the serialized public-key, the size of *sigma*, the number of polynomial coefficients in $\boldsymbol{B}$ and the size in bits of each polynomial coefficient.

unpack_ct

Function that de-serializes the components $\boldsymbol{U}$ and $\boldsymbol{v}$ of Round5's ciphertext *ct*. The function is defined as follows and uses *unpack* internally.

$(\boldsymbol{U}, \boldsymbol{v}) = unpack\_ct(ct, U\_elem, U\_elem\_size, v\_elem, v\_elem\_size)$

where the input parameters are the serialized ciphertext, the number of polynomial coefficients in $\boldsymbol{U}$, the size in bits of each polynomial coefficient, the number of polynomial coefficients in $\boldsymbol{v}$, and the size in bits of each vector coefficient.

verify

Compares two byte strings $s1$ and $s2$ of equal length l and outputs bit 0 if they are equal.

$$c = \text{verify}(s1, s2, l)$$

conditional _constant _time_memcpy

> Function that copies byte string a of length a_bytes starting at memory address mem if condition $cond$ is true.

$$conditional\_constant\_time\_memcpy(mem, a, a\_bytes, cond)$$

1.11.7 Implementation of r5_cpa_pke

Round5's INDCPA public-key encryption is specified in Algorithms 1, 2, and 3. Their implementation is described in Algorithms 13, 14, and 15.

Algorithm 13: r5_cpa_pke_keygen

---

output: $pk : kappa\_bytes + \lceil n\_bar \cdot d/n \cdot n \cdot p\_bits/8 \rceil$ byte string.
     : $sk : kappa\_bytes$ byte string.

1 $sigma = randombytes(kappa\_bytes)$

2 $\boldsymbol{A} = create\_A(sigma)$

3 $sk = randombytes(kappa\_bytes)$

4 $\boldsymbol{S\_T} = create\_S\_T(sk)$

5 $\boldsymbol{S} = transpose\_matrix(S\_T, n\_bar, d/n, n)$

6 $\boldsymbol{B} = mult\_matrix(\boldsymbol{A}, d/n, d/n, \boldsymbol{S}, d/n, n\_bar, n, q, Phi)$

7 $\boldsymbol{B} = round\_matrix(\boldsymbol{B}, d/n \cdot n\_bar, n, q\_bits, p\_bits, h1)$

8 $pk = pack\_pk(sigma, kappa\_bytes, b\_bits, d/n \cdot n\_bar \cdot n, p\_bits)$

9 **return** pk, sk

Algorithm 14: r5_cpa_pke_encrypt

---

input: $pk : kappa\_bytes + \lceil n\_bar \cdot d/n \cdot n \cdot p\_bits/8 \rceil$ byte string.

input: $m : kappa\_bytes$ byte string that is encrypted.

input: $\rho : kappa\_bytes$ byte string.

output: $ct : (\lceil n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil$ encrypted byte string.

1 $sigma, \boldsymbol{B} = unpack\_pk(pk, kappa\_bytes, d/n \cdot n\_bar \cdot n, p\_bits)$

2 $\boldsymbol{A} = create\_A(sigma)$

3 $\boldsymbol{R\_T} = create\_R\_T(\rho)$

4 $\boldsymbol{A\_T} = transpose\_matrix(\boldsymbol{A}, d/n, d/n, n)$

5 $\boldsymbol{R} = transpose\_matrix(\boldsymbol{R\_T}, m\_bar, d/n, n)$

6 $\boldsymbol{U} = mult\_matrix(\boldsymbol{A\_T}, d/n, d/n, \boldsymbol{R}, d/n, m\_bar, n, q, Phi)$

7 $\boldsymbol{U} = round\_matrix(\boldsymbol{U}, d/n \cdot m\_bar, n, q\_bits, p\_bits, h2)$

8 $\boldsymbol{B\_T} = transpose\_matrix(\boldsymbol{B}, d/n, n\_bar, n)$

9 $\boldsymbol{X} = mult\_matrix(\boldsymbol{B\_T}, n\_bar, d/n, \boldsymbol{R}, d/n, m\_bar, n, p, Xi)$

10 $\boldsymbol{x} = round\_matrix(\boldsymbol{sample\_mu(X)}, mu, 1, p\_bits, t\_bits, h2)$

11 $m1 = xef\_compute(m)$

12 $\boldsymbol{v} = add\_msg(mu, \boldsymbol{x}, m1, b\_bits, t\_bits)$

13 $ct = pack\_ct(\boldsymbol{U}, d/n \cdot m\_bar \cdot n, p\_bits, \boldsymbol{v}, mu, t\_bits)$

14 **return** ct

Algorithm 15: r5_cpa_pke_decrypt

---

input: $ct : (\lceil n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil$ byte string.
     : $sk : kappa\_bytes$ byte string.
output: $m : kappa\_bytes$ byte string that has been decrypted.

1 $\boldsymbol{S\_T} = create\_S\_T(sk)$
2 $\boldsymbol{U}, \boldsymbol{v} = unpack\_ct(ct, d/n \cdot m\_bar \cdot n, p\_bits, mu, t\_bits)$
3 $\boldsymbol{v} = decompress\_matrix(\boldsymbol{v}, mu, 1, p\_bits, t\_bits)$
4 $\boldsymbol{X\_prime} = mult\_matrix(\boldsymbol{S\_T}, n\_bar, d/n, \boldsymbol{U}, d/n, m\_bar, n, p, Xi)$
5 $\boldsymbol{m2} = diff\_msg(mu, \boldsymbol{v}, \boldsymbol{sample\_mu}(\boldsymbol{X\_prime}), p)$
6 $\boldsymbol{m2} = round\_matrix(\boldsymbol{m2}, mu, 1, p\_bits, b\_bits, h3)$
7 $m1 = pack(\boldsymbol{m2}, mu, b\_bits)$
8 $m1 = xef\_compute(m1, kappa\_bytes, f)$
9 return $m = xef\_fixerr(m1, kappa\_bytes, f)$

1.11.8 Implementation of r5_cpa_kem

Round5's main building block is an INDCPA KEM are specified in Algorithms 4, 5, and 6. They are described from an implementation perspective in Algorithms 16, 17, and 18.

Algorithm 16: r5_cpa_kem_keygen

---

output: $pk : \lceil kappa\_bytes + n\_bar \cdot d/n \cdot n \cdot p\_bits/8 \rceil$ byte string.
     : $sk : kappa\_bytes$ byte string.

1 $(pk, sk) = r5\_cpa\_pke\_keygen$
2 return pk, sk

Algorithm 17: r5_cpa_kem_encapsulate

---

input: $pk : \lceil kappa\_bytes + n\_bar \cdot d/n \cdot n \cdot p\_bits/8 \rceil$ byte string.
output: $ct : (\lceil n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil$ byte string.
output: $k : kappa\_bytes$ byte string.

1 $m = randombytes(kappa\_bytes)$
2 $rho = randombytes(kappa\_bytes)$
3 $ct = r5\_cpa\_pke\_encrypt(pk, m, rho)$
4 $k = hash(kappa\_bytes, m||ct, "", kappa\_bytes + \lceil n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil, "", 0)$
5 return k

Algorithm 18: r5_cpa_pke_decapsulate

---

input: $sk : kappa\_bytes$ byte string.
input: $ct : (\lceil n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil$ byte string.
output: $k : kappa\_bytes$ byte string.

1 $m = r5\_cpa\_pke\_decrypt(sk, ct)$
2 $k = hash(kappa\_bytes, m||ct, kappa\_bytes + \lceil n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil, "", 0)$
3 return k

1.11.9 Implementation of r5_cca_kem

Round5 relies on an INDCCA KEM specified in Algorithms 7, 8, and 9. They are described from an implementation perspective in Algorithms 19, 20, and 21.

Algorithm 19: r5_cca_kem_keygen

---

output: $pk : \lceil kappa\_bytes + n\_bar \cdot d/n \cdot n \cdot p\_bits/8 \rceil$ byte string.
    : $sk : \lceil 3 \cdot kappa\_bytes + n\_bar \cdot d/n \cdot n \cdot p\_bits/8 \rceil$ byte string.

1 $(pk, sk) = r5\_cpa\_pke\_keygen$

2 $y = randombytes(kappa\_bytes)$

3 $sk = sk||y||pk$

4 return (pk, sk)

Algorithm 20: r5_cca_kem_encapsulate

---

input: $pk : \lceil 3 \cdot kappa\_bytes + (n\_bar \cdot d/n \cdot n \cdot p\_bits)/8 \rceil$ byte string.
output: $ct : \lceil kappa\_bytes + (n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil$ byte string.
output: $k : kappa\_bytes$ byte string.

1 $m = randombytes(kappa\_bytes)$

2 $L||g||rho =$
    $hash(3 \cdot kappa\_bytes, m||pk, \lceil 2 \cdot kappa\_bytes + (n\_bar \cdot d/n \cdot n \cdot p\_bits)/8 \rceil, "", 0)$

3 $(\boldsymbol{U}, \boldsymbol{v}) = r5\_cpa\_pke\_encrypt(pk, m, rho)$

4 $ct = U||v||g$

5 $k = hash(kappa\_bytes, L||ct, kappa\_bytes + \lceil (n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil, "", 0)$

6 return ct, k

Algorithm 21: r5_cca_pke_decapsulate

---

input: $sk : \lceil 3 \cdot kappa\_bytes + n\_bar \cdot d/n \cdot n \cdot p\_bits/8 \rceil$ byte string.
input: $ct : \lceil kappa\_bytes + (n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil$ byte string.
output: $k : kappa\_bytes$ byte string.

1 $m' = r5\_cpa\_pke\_decrypt(sk, ct)$

2 $L\_prime||g\_prime||rho\_prime =$
    $hash(3 \cdot kappa\_bytes, m'||pk, \lceil 2 \cdot kappa\_bytes + (n\_bar \cdot d/n \cdot n \cdot p\_bits)/8 \rceil, "", 0)$

3 $(\boldsymbol{U\_prime}, \boldsymbol{v\_prime}) = r5\_cpa\_pke\_encrypt(pk, m\_prime, rho\_prime)$

4 $input = L\_prime||ct$

5 $fail = verify(ct, ct\_prime, \lceil kappa\_bytes + (n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil)$

6 $conditional\_constant\_time\_memcpy(hash_input, y, kappa\_bytes, fail)$

7 $k = hash(kappa\_bytes, hash\_input, 2 \cdot kappa\_bytes + \lceil n\_bar \cdot d/n \cdot n \cdot p\_bits + mu \cdot t\_bits)/8 \rceil, "", 0)$

8 return k

A Formal security of Round5

This section contains details on the formal security of Round5, its algorithms and building blocks. It is organized as follows: Section A.2 introduces notions of security based on indistinguishablity of ciphertexts or encapsulated keys. Section A.3 gives results (existing and new) on the hardness of our underlying problem, or rather its two specific instances we use – the Learning with Rounding problem with sparse ternary secrets ($\texttt{LWR}_{\text{spt}}$) and the Ring Learning with Rounding problem with sparse ternary secrets ($\texttt{RLWR}_{\text{spt}}$). Sections A.6 and A.7 contain the first main result of this section – a proof of IND-CPA security for r5_cpa_kem (Theorem A.6.0.1), and a proof of IND-CCA security of r5_cca_pke (Theorems A.7.0.1 and A.7.0.2), assuming the hardness of the above problems. Finally, Section A.8 and Theorem A.8.0.1 contain the second main result of this reduction: a proof of the hardness for $\texttt{LWR}_{\text{spt}}$, the underlying problem of our schemes for the non-ring case (i.e., for $n = 1$) in the form of a polynomial-time reduction to it from the Learning with Errors (LWE) problem with secrets uniformly chosen from $\mathbb{Z}_q^d$ and errors drawn according to a Gaussian distribution.

A.1 Deterministic generation of $\boldsymbol{A}$

The General Learning with Rounding (GLWR) public parameter $\boldsymbol{A}$ in Round5 is generated using the function $f_{d,n}^{(\tau)}$ from a short random seed (see Section 1.4.2). The core component in $f_{d,n}^{(\tau)}$ responsible for deterministically expanding this short random seed into a longer random sequence is either AES(128 or 256) [44] or SHAKE(128 or 256) [45]. In order to relate Round5's security to the hardness of the GLWR problem, we reuse Naehrig *et al.*'s argument in [75] to argue that we can replace a uniformly sampled matrix $\boldsymbol{A} \in \mathcal{R}_{n,q}^{d/n \times d/n}$ with matrices sampled according to Round5's key-generation algorithm, for both of the above two algorithms, while considering a realistic adversary with access to the seed. The proof for both the cases of AES and SHAKE proceeds by using the notion of indifferentiability [71, 37, Def. 3], in exactly the same manner as in [75, Sec. 5.1.4].

In case of AES, the proof holds directly for the instantiation $f_{d,n}^{(0)}$, and also for $f_{d,n}^{(1)}$ when the function permutes complete AES blocks. We explain the intuition behind the proof. Let $\mathcal{F}$ denote an "ideal domain expansion" primitive that expands a short random seed, block-wise, into a larger sequence, such that each block is unique and also sampled uniformly at random. In our security reductions, the GLWR public parameter $\boldsymbol{A}$ is generated by querying the GLWR oracle. It can be shown that marginally increasing the number of calls to the GLWR oracle makes it possible to construct a GLWR matrix $\boldsymbol{A}$ that fits (with high probability) the output distribution of $\mathcal{F}$, without deteriorating the problem's hardness [75, Sec. 5.1.4]. Next, we consider a construction $\mathcal{C}^{\mathcal{G}}$ in the Ideal Cipher model implementing $\mathcal{F}$ as AES (as in Round5). It can be shown that $\mathcal{C}^{\mathcal{G}}$ is indifferentiable from $\mathcal{F}$ [75, Sec. 5.1.4]. This therefore allows us to replace

the uniformly random sampling of $\boldsymbol{A}$ in the GLWR problem with one generated as in Round5's $f_{d,n}^{(\tau)}$ without affecting security. By definition, this is directly possible if $\tau = 0$; it also holds for $\tau = 1$ if $f_{d,n}^{(1)}$ permutes complete AES blocks.

Next, we explain the intuition behind the proof when SHAKE is used in $f_n^{(0)|}$ or $f_n^{(1)}$. In the random oracle model, SHAKE is an ideal XOF [43]. It can be shown that [75, Sec. 5.1.4] SHAKE can be modeled as an ideal hash function used to expand a seed into each row of the matrix $\boldsymbol{A}$, each step being independent, thereby expanding the uniformly random seed into a larger uniformly random matrix. This construction implements the ideal functionality $\mathcal{F}$ perfectly, completing the proof. We refer to [75, Sec. 5.1.4] for details.

A.2   Security Definitions

Security requirements for public-key encryption and key-encapsulation schemes are based on whether static or ephemeral keys are used. (Static) public-key encryption (PKE) schemes, and nominally ephemeral key-encapsulation mechanisms that allow key caching that provide security against adaptive chosen ciphertext attack (corresponding to IND-CCA2 security) are considered to be sufficiently secure [76, Section 4.A.2]. On the other hand, a purely ephemeral key encapsulation mechanism (KEM) that provides semantic security against chosen plaintext attack, i.e., IND-CPA security, is considered to be sufficiently secure [76, Section 4.A.3].

Definition A.2.0.1 (Distinguishing advantage). *Let $\mathcal{A}$ be a randomized algorithm that takes as input elements from a set X with output in $\{0, 1\}$. Let D_1 and D_2 two probability distributions on X. The advantage of $\mathcal{A}$ for distinguishing between D_1 and D_2, denoted by $Adv_{D_1, D_2}(\mathcal{A})$, is defined as*

$$Adv_{D_1, D_2}(\mathcal{A}) = |Pr[\mathcal{A}(x) = 1 \mid x \leftarrow D_1] - Pr[\mathcal{A}(x) = 1 \mid x \leftarrow D_2]\,|$$

Table 20: IND-CPA game for PKE

1. $(pk, sk) = \texttt{KeyGen}(\lambda)$.
2. $b \xleftarrow{\$} \{0, 1\}$
3. $(m_0, m_1, st) = \mathcal{A}(pk)$ such that $|m_0| = |m_1|$.
4. $c = \texttt{Enc}(pk, m_b)$
5. $b' = \mathcal{A}(pk, c, st)$
6. **return** $[b' = b]$

Definition A.2.0.2 (IND-CPA Secure PKE). *Let PKE=(Keygen, Enc, Dec) be a public key encryption scheme with message space $\mathcal{M}$. Let λ be a security parameter. The IND-CPA game is defined in Table 20, and the IND-CPA*

advantage of an adversary $\mathcal{A}$ against PKE is defined as

$$Adv_{PKE}^{\textit{IND-CPA}}(\mathcal{A}) \ = \ \mid Pr[\textit{IND-CPA}^{\mathcal{A}} \Rightarrow 1] - \frac{1}{2} \mid .$$

Definition A.2.0.3 (IND-CCA Secure PKE)**.** *Let PKE=(Keygen, Enc, Dec) be a public key encryption scheme with message space $\mathcal{M}$. Let λ be a security parameter. Let $\mathcal{A}$ be an adversary against PKE. The IND-CCA game is defined as Table 20, with the addition that $\mathcal{A}$ has access to a decryption oracle $Dec(\cdot) = Dec(sk, \cdot)$ that returns $m' = Dec(sk, Enc(pk, m'))$, and the restriction that $\mathcal{A}$ cannot query $Dec(\cdot)$ with the challenge c. The IND-CCA advantage of $\mathcal{A}$ against PKE is defined as*

$$Adv_{PKE}^{\textit{IND-CCA}}(\mathcal{A}^{Dec(\cdot)}) \ = \ \mid Pr[\textit{IND-CCA}^{\mathcal{A}} \Rightarrow 1] - \frac{1}{2} \mid .$$

Table 21: IND-CPA game for KEM

---

1. $(pk, sk) = \texttt{KeyGen}(\lambda)$.
2. $b \xleftarrow{\$} \{0, 1\}$
3. $(c, K_0) = \texttt{Encaps}(pk)$
4. $K_1 \xleftarrow{\$} \mathcal{K}$
5. $b' = \mathcal{A}(pk, c, K_b)$
6. **return** $[b' = b]$

Definition A.2.0.4 (IND-CPA Secure KEM)**.** *Let KEM=(Keygen, Encaps, Decaps) be a key encapsulation mechanism with key space $\mathcal{K}$. Let λ be a security parameter. The IND-CPA game is defined in Table 21, and the IND-CPA advantage of an adversary $\mathcal{A}$ against KEM is defined as*

$$Adv_{KEM}^{\textit{IND-CPA}}(\mathcal{A}) \ = \ \mid Pr[\textit{IND-CPA}^{\mathcal{A}} \Rightarrow 1] - \frac{1}{2} \mid .$$

Definition A.2.0.5 (IND-CCA Secure KEM)**.** *Let KEM=(Keygen, Encaps, Decaps) be a key encapsulation mechanism with key space $\mathcal{K}$. Let λ be a security parameter. Let $\mathcal{A}$ be an adversary against KEM. The IND-CCA game is defined in Table 21, with the addition that $\mathcal{A}$ has access to a decapsulation oracle $Decaps(\cdot) = Decaps(sk, \cdot)$ that returns $K' = Decaps(sk, c')$ where $(c', K') = Encaps(pk)$, and the restriction that $\mathcal{A}$ cannot query $Decaps(\cdot)$ with the challenge c. The IND-CCA advantage of $\mathcal{A}$ against KEM is defined as*

$$Adv_{KEM}^{\textit{IND-CCA}}(\mathcal{A}^{Decaps(\cdot)}) \ = \ \mid Pr[\textit{IND-CCA}^{\mathcal{A}} \Rightarrow 1] - \frac{1}{2} \mid .$$

In the random oracle model [21] (both for IND-CPA and IND-CCA games), the adversary is given access to a random oracle H that it can query up to a

polynomial number q_H of times. In a post-quantum setting, it can be assumed that the adversary has access to a quantum accessible random oracle H_Q [24] that can be queried up to q_{H_Q} times on arbitrary superpositions of input strings.

A.3   Hardness Assumption (Underlying Problem)

In this section, we detail the underlying problem on whose hardness the security of our schemes are established. Depending on whether the system parameter n is chosen to be 1 or d (see Section 1.4.2), the proposed public-key encryption and key-encapsulation mechanism are instantiated either as non-ring (LWR) based or ring (RLWR) based schemes. The security of the proposals are therefore based on the hardness of the Decision-General Learning with Rounding problem with sparse-ternary secrets, i.e., $\mathtt{dGLWR}_{\mathrm{spt}}$ (see Section 1.3). We first recall hardness results for LWR before introducing our own results on the hardness of $\mathtt{dLWR}_{\mathrm{spt}}$. We then recall hardness results of RLWR.

Provable Security in the non-ring case: The hardness of the LWR problem has been studied in [18, 9, 23, 17] and established based on the hardness of the Learning with Errors (LWE) problem [83]. The most recent work are two independent reductions to LWR from LWE: the first, due to Bai et al. [17, Theorem 6.4] preserves the dimension n between the two problems but decreases the number of LWR samples that may be queried by an attacker by a factor (p/q); the second due to Bogdanov et al. [23, Theorem 3] preserves the number of samples between the two problems but increases the LWR dimension by a factor $\log q$. We follow the approach of Bai since it results in a smaller LWR dimension, leading to smaller bandwidth requirements and better performance.

Bai et al.'s reduction [17, Theorem 6.4] is an essential component that we use in Section A.8 to prove a reduction from $\mathtt{dLWE}_{n,m',q,D_\alpha}\left(\mathcal{U}\left(\mathbb{Z}_q\right)\right)$ to $\mathtt{dLWR}_{n,m,q,p}\left(\mathcal{U}(\mathcal{H}_n\left(h\right))\right)$, i.e., to $\mathtt{dLWR}_{\mathrm{spt}}$.

Provable Security in the ring case: Next, we recall hardness results for the ring case, i.e., for Decision-RLWR [18]. To the best of our knowledge, the only existing result on the hardness of Decision-RLWR is due to [18, Theorem 3.2], who show that Decision-RLWR is at least as hard as Decision-RLWE as long as the underlying ring and secret distribution remain the *same* for the two problems, the RLWE noise is sampled from *any* (balanced) distribution in $\{-B, \ldots, B\}$, and q is super-polynomial in n, i.e., $q \geq pBn^{\omega(1)}$. The last condition may be too restrictive for practical schemes. Hence, although [18, Theorem 3.2] is relevant for the provable (IND) security of our schemes' ring-based instantiations, it remains to be seen whether the above reduction can be improved to be made practical.

A.4   IND-CPA Security of r5_cpa_pke

In this section it is shown that the public-key encryption scheme r5_cpa_pke is IND-CPA secure, based on the hardness of the decision GLWR problem

with sparse-ternary secrets. In Section A.7 this result will be used to show that r5_cca_pke is IND-CCA Secure. The proof is restricted to the case that $\xi(x) = \Phi_{n+1}(x)$. Referring to Table 20, we take for KeyGen the function r5_cpa_pke_keygen, and for Enc the function r5_cpa_pke_encrypt(pk, m, ρ) with uniform choice for ρ. That is,

$$\rho \xleftarrow{\$} \{0,1\}^{\kappa}; \text{Enc}(pk, m) = \text{r5\_cpa\_pke\_encrypt}(pk, m, \rho).$$

Let n, m, p, q, d, h be positive integers with $n \in \{1, d\}$. The hard problem underlying the security of our schemes is decision-GLWR with sparse-ternary secrets (see Section 1.3). For the above parameters, we define the GLWR oracle $O_{m,\chi_S,\boldsymbol{s}}$ for a secret distribution χ_S that returns m GLWR samples as follows:

$$O_{m,\chi_S,\boldsymbol{s}} : \boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{m \times d/n}, \ \boldsymbol{s} \leftarrow \chi_S; \ \text{return} \ (\boldsymbol{A}, \text{R}_{q \to p}(\boldsymbol{A}\boldsymbol{s})) \tag{35}$$

The decision-GLWR problem with sparse-ternary secrets is to distinguish between the distributions $\mathcal{U}(\mathcal{R}_{n,q}^{m \times d/n}) \times \mathcal{U}(\mathcal{R}_{n,p})$ and $O_{m,\chi_S,\boldsymbol{s}}$, with $\boldsymbol{s}$ common to all samples and $\chi_S := \mathcal{U}(\mathcal{H}_{n,d/n}(h))$. For an adversary $\mathcal{A}$, we define

$$\text{Adv}_{d,n,m,q,p}^{\texttt{dGLWR}_{\texttt{spt}}}(\mathcal{A}) =$$

$$\left| \Pr\left[\mathcal{A}(\boldsymbol{A}, \boldsymbol{b}) = 1 \mid (\boldsymbol{A}, \boldsymbol{b}) \xleftarrow{\$} O_{m,\chi_S,\boldsymbol{s}} \right] - \Pr\left[\mathcal{A}(\boldsymbol{A}, \boldsymbol{b}) = 1 \mid \boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{m \times d/n}, \boldsymbol{b} \xleftarrow{\$} \mathcal{R}_{n,p}^{m} \right] \right|$$

For an extended form of the decision-GLWR problem with the secret in form of a matrix consisting of $\overline{n}$ independent secret vectors, we define a similar oracle $O_{m,\chi_S,\overline{n},\boldsymbol{S}}$ as follows:

$$O_{m,\chi_S,\overline{n},\boldsymbol{S}} : \boldsymbol{A} \xleftarrow{\$} \mathcal{U}\left(\mathcal{R}_{n,q}^{m \times d/n} \right), \ \boldsymbol{S} \leftarrow (\chi_S)^{\overline{n}}; \ \text{return} \ (\boldsymbol{A}, \text{R}_{q \to p}(\boldsymbol{A}\boldsymbol{S})) \tag{36}$$

The advantage of an adversary for this extended form of the decision-GLWR problem is defined in a similar manner as above.

The following theorem shows that r5_cpa_pke is IND-CPA secure assuming the hardness of decision-GLWR with sparse-ternary secrets.

Theorem A.4.0.1. *Let p, q, t be integers such that $t|p|q$, and let $z = \max(p, tq/p)$. Furthermore, assume that $\xi(x) = \Phi_{n+1}(x)$. If $f_{d,n}^{(\tau)}$, f_S and f_R induce distributions indistinguishable from uniform, then r5_cpa_pke is IND-CPA secure under the hardness assumption of the Decision-GLWR problem with sparse-ternary secrets. More precisely, for every IND-CPA adversary $\mathcal{A}$, if $Adv_{CPA\text{-}PKE}^{IND\text{-}CPA}(\mathcal{A})$ is the advantage in winning the IND-CPA game, then there exist distinguishers $\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}, \mathcal{F}$ such that*

$$\begin{aligned} Adv_{CPA\text{-}PKE}^{IND\text{-}CPA}(\mathcal{A}) \leq &\ Adv_{\mathcal{U}(\mathcal{R}_{n,q}^{d/n \times d/n}), F_\tau}(\mathcal{B}) + Adv_{G_S, \chi_S^{\overline{n}}}(\mathcal{C}) + Adv_{G_R, \chi_S^{\overline{m}}}(\mathcal{D}) \\ &+ \overline{n} \cdot Adv_{d,n,d/n,q,p}^{\texttt{dGLWR}_{spt}}(\mathcal{E}) + Adv_{d,n,d/n+\overline{n},q,z}^{\texttt{dGLWR}_{spt}}(\mathcal{F}) \end{aligned} \tag{37}$$

In this equation, F_τ is the distribution of $f_{d,n}^{(\tau)}$ with $\sigma \xleftarrow{\$} \{0,1\}^{\kappa}$, G_S is the distribution of $f_S(s)$ with $s \xleftarrow{\$} \{0,1\}^{\kappa}$ and G_R is the distribution of $f_R(\rho)$ with

Table 22: IND-CPA games for r5_cpa_pke: games G_0 and G_1

Game G_0	Game G_1
1. $\sigma \xleftarrow{\$} \{0,1\}^\kappa$; $\boldsymbol{A} = f_{d,n}^{(\tau)}(\sigma)$	1. $\sigma \xleftarrow{\$} \{0,1\}^\kappa$; $\boldsymbol{A} = f_{d,n}^{(\tau)}(\sigma)$
2. $s \xleftarrow{\$} \{0,1\}^\kappa$; $\boldsymbol{S} = f_S(s)$	2. $s \xleftarrow{\$} \{0,1\}^\kappa$; $\boldsymbol{S} = f_S(s)$
3. $\boldsymbol{B} = \mathrm{R}_{q \to p}(\langle \boldsymbol{AS} \rangle_{\Phi_{n+1}})$	3. $\boldsymbol{B} = \mathrm{R}_{q \to p}(\langle \boldsymbol{AS} \rangle_{\Phi_{n+1}})$
4. Choose $\beta \xleftarrow{\$} \{0,1\}$.	4. Choose $\beta \xleftarrow{\$} \{0,1\}$.
5. $(m_0, m_1, st) = \mathcal{A}(\boldsymbol{A}, \boldsymbol{B})$	5. $(m_0, m_1, st) = \mathcal{A}(\boldsymbol{A}, \boldsymbol{B})$
6. $\rho \xleftarrow{\$} \{0,1\}^\kappa$; $\boldsymbol{R} = f_R(\rho)$	6. $\rho \xleftarrow{\$} \{0,1\}^\kappa$; $\boldsymbol{R} = f_R(\rho)$
7. $\boldsymbol{U} = \mathrm{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$	7. $\boldsymbol{U} = \mathrm{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$
8. $\boldsymbol{v} = \langle \mathrm{Sample}_\mu \left(\mathrm{R}_{p \to t, h_2}(\langle \boldsymbol{B}^T \boldsymbol{R} \rangle_\xi) \right) + \frac{t}{b} ECC\_Enc_{\kappa,f}(m_\beta) \rangle_t$	8. $\boldsymbol{V} = \langle \left(\mathrm{R}_{p \to t, h_2}(\langle \boldsymbol{B}^T \boldsymbol{R} \rangle_\xi) \right) + \frac{t}{b} \boldsymbol{M}_\beta \rangle_t$
9. $\beta' = \mathcal{A}\left((\boldsymbol{A}, \boldsymbol{B}), (\boldsymbol{U}, \boldsymbol{v}), st \right)$	9. $\beta' = \mathcal{A}\left((\boldsymbol{A}, \boldsymbol{B}), (\boldsymbol{U}, \mathrm{Sample}_\mu(\boldsymbol{V})), st \right)$
10. **return**$[(\beta' = \beta)]$.	10. **return**$[(\beta' = \beta)]$.

$\rho \xleftarrow{\$} \{0,1\}^\kappa$. *Moreover, $Adv_{d,n,m,q_1,q_2}^{dGLWR_{spt}}(\mathcal{Z})$ is the advantage of adversary $\mathcal{Z}$ in distinguishing m GLWR samples (with sparse-ternary secrets) from uniform, with the GLWR problem defined for the parameters d, n, q_1, q_2.*

Theorem A.4.0.1 via a sequence of IND-CPA games shown in Tables 22 to 25, following the methodology of Peikert et al. in [77, Lemma 4.1] and that of [26, Theorem 3.3]. Steps for combining samples using rounding from q to p and samples using rounding from p to t are due to [40]. Game G_0 is the actual CPA-PKE game. For convenience, the steps of r5_cpa_pke_encrypt(pk, m, ρ) are written out explicitly. Moreover, we write χ_S for the uniform distribution on $\mathcal{H}_{n,d/n}(h)$.

In Game G_1, $\boldsymbol{M}_\beta$ is the matrix that has zero coefficients in all positions not picked up by Sample_μ and satisfies $ECC\_Enc_{\kappa,f}(m_\beta) = \mathrm{Sample}_\mu(\boldsymbol{M}_\beta)$. Clearly,

$$\Pr(S_0) = \Pr(S_1) \tag{38}$$

Games G_1 and G_2 only differ in the generation of $\boldsymbol{A}$. A distinguisher between $\mathcal{U}(\mathcal{R}_{n,q}^{d/n \times d/n})$ and F_τ, defined as the distribution of $f_{d,n}^\tau(\sigma)$ with $\sigma \xleftarrow{\$} \{0,1\}^\kappa$, can be constructed as follows. On input $\boldsymbol{A} \in \mathcal{R}_{n,q}^{d/n \times d/n}$, perform steps 2-10 of game G_1. The output b' is distributed as in game G_1 if $\boldsymbol{A}$ is distributed according to F_τ, and is distributed as in game G_2 if $\boldsymbol{A}$ is distributed uniformly. We conclude that there is a distinghuisher $\mathcal{B}$ such that

$$\mathrm{Adv}_{\mathcal{U}(\mathcal{R}_{n,q}^{d/n \times d/n}), F_\tau}(\mathcal{B}) = |\Pr(S_1) - \Pr(S_2)|. \tag{39}$$

Games G_2 and G_3 only differ in the generation of the secret matrix $\boldsymbol{S}$. We denote the distribution of $f_S(s)$ with $s \xleftarrow{\$} \{0,1\}^\kappa$ by G_S. Similarly to the

Table 23: IND-CPA games for r5_cpa_pke: games G_2 and G_3

Game G_2	Game G_3
1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$	1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$
2. $s \xleftarrow{\$} \{0,1\}^{\kappa}$; $\boldsymbol{S} = f_S(s)$	2. $\boldsymbol{S} \xleftarrow{\$} \chi_S^{\overline{n}}$
3. $\boldsymbol{B} = \mathrm{R}_{q \to p}(\langle \boldsymbol{AS} \rangle_{\Phi_{n+1}})$	3. $\boldsymbol{B} = \mathrm{R}_{q \to p}(\langle \boldsymbol{AS} \rangle_{\Phi_{n+1}})$
4. Choose $\beta \xleftarrow{\$} \{0,1\}$.	4. Choose $\beta \xleftarrow{\$} \{0,1\}$.
5. $(m_0, m_1, st) = \mathcal{A}(\boldsymbol{A}, \boldsymbol{B})$.	5. $(m_0, m_1, st) = \mathcal{A}(\boldsymbol{A}, \boldsymbol{B})$.
6. $\rho \xleftarrow{\$} \{0,1\}^{\kappa}$; $\boldsymbol{R} = f_R(\rho)$	6. $\rho \xleftarrow{\$} \{0,1\}^{\kappa}$; $\boldsymbol{R} = f_R(\rho)$
7. $\boldsymbol{U} = \mathrm{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\phi_{n+1}})$	7. $\boldsymbol{U} = \mathrm{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\phi_{n+1}})$
8. $\boldsymbol{V} = \langle (\mathrm{R}_{p \to t, h_2}(\langle \boldsymbol{B}^T \boldsymbol{R} \rangle_{\xi})) + \frac{t}{b} \boldsymbol{M_b} \rangle_t$	8. $\boldsymbol{V} = \langle (\mathrm{R}_{p \to t, h_2}(\langle \boldsymbol{B}^T \boldsymbol{R} \rangle_{\xi})) + \frac{t}{b} \boldsymbol{M_b} \rangle_t$
9. $\beta' = \mathcal{A}((\boldsymbol{A}, \boldsymbol{B}), (\boldsymbol{U}, \mathrm{Sample}_{\mu}(\boldsymbol{V})), st)$	9. $\beta' = \mathcal{A}((\boldsymbol{A}, \boldsymbol{B}), (\boldsymbol{U}, \mathrm{Sample}_{\mu}(\boldsymbol{V})), st)$
10. **return**$[(\beta' = \beta)]$.	10. **return**$[(\beta' = \beta)]$.

reasoning for Games G_1 and G_2, there is a distinguisher $\mathcal{C}$ such that

$$\mathrm{Adv}_{G_S, \chi_S^{\overline{n}}}(\mathcal{C}) = |\Pr(S_2) - \Pr(S_3)|. \tag{40}$$

Games G_3 and G_4 only differ in the generation of $\boldsymbol{B}$. Consider the following algorithm. On input $(\boldsymbol{A}, \boldsymbol{B}) \in \mathcal{R}_{n,q}^{d/n \times d/n} \times \mathcal{R}_{n,p}^{d/n \times \overline{n}}$, run steps 3-10 from game G_3. As $(\mathbf{A}, \mathbf{B})$ is distributed like $O_{m, \chi_S, \overline{n}, \boldsymbol{S}}$ in game G_3 and uniformly in game G_4, we conclude that there is a distinguisher $\mathcal{X}$ such that

$$\mathrm{Adv}_{O_{m, \chi_S, \overline{n}, \boldsymbol{S}, \mathcal{U}(\mathcal{R}_{n,q}^{d/n \times d/n} \times \mathcal{R}_{n,p}^{d/n \times \overline{n}})}}(\mathcal{X}) = |\Pr(S_3) - \Pr(S_4))| \tag{41}$$

By using a standard hybrid argument, we infer that there is a distinguisher $\mathcal{E}$ such that

$$\mathrm{Adv}_{O_{m, \chi_S, \overline{n}, \boldsymbol{S}, \mathcal{U}(\mathcal{R}_{n,q}^{d/n \times d/n} \times \mathcal{R}_{n,p}^{d/n \times \overline{n}})}}(\mathcal{X}) \leq \overline{n} \cdot \mathrm{Adv}_{O_{m, \chi_S, \overline{1}, \boldsymbol{s}, \mathcal{U}(\mathcal{R}_{n,q}^{d/n \times d/n} \times \mathcal{R}_{n,p}^{d/n \times 1})}}(\mathcal{E}) \tag{42}$$

Games G_4 and G_5 only differ in the generation of the secret matrix $\boldsymbol{R}$. We denote the distribution of $f_R(\rho)$ with $\rho \xleftarrow{\$} \{0,1\}^{\kappa}$ by G_R. Similarly to the reasoning for games G_1 and G_2, there is a distinguisher $\mathcal{D}$ such that

$$\mathrm{Adv}_{G_R, \chi_S^{\overline{m}}}(\mathcal{D}) = |\Pr(S_4) - \Pr(S_5)|. \tag{43}$$

Now consider Games G_5 and G_6. As p divides q, $\langle \boldsymbol{B}_q \rangle_p$ is uniformly distributed. By definition of the rounding function,

$$\boldsymbol{V'} \equiv \lfloor \frac{t}{p}(\langle \boldsymbol{B_q}^T \boldsymbol{R} \rangle_{\xi} + h_2 \boldsymbol{J}) \rceil + \frac{t}{b} \boldsymbol{M_b} \pmod{\frac{tq}{p}}.$$

As p divides q and $\boldsymbol{B}_q \equiv \langle \boldsymbol{B_q} \rangle_p \pmod{p}$,

$$\boldsymbol{V''} \equiv \lfloor \frac{t}{p}(\langle \boldsymbol{B_q}^T \rangle_p \boldsymbol{R} \rangle_{\xi} + h_2 \boldsymbol{J}) \rceil + \frac{t}{b} \boldsymbol{M_b} \pmod{t}.$$

Table 24: IND-CPA games for r5_cpa_pke: games G_4 and G_5

Game G_4	Game G_5
1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$	1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$
2. -	2. -
3. $\boldsymbol{B} \xleftarrow{\$} \mathcal{R}_{n,p}^{d/n \times \overline{n}}$	3. $\boldsymbol{B} \xleftarrow{\$} \mathcal{R}_{n,p}^{d/n \times \overline{n}}$
4. Choose $\beta \xleftarrow{\$} \{0,1\}$.	4. Choose $\beta \xleftarrow{\$} \{0,1\}$.
5. $(m_0, m_1, st) = \mathcal{A}(\boldsymbol{A}, \boldsymbol{B})$.	5. $(m_0, m_1, st) = \mathcal{A}(\boldsymbol{A}, \langle(\boldsymbol{B_q})_p\rangle)$.
6. $\rho \xleftarrow{\$} \{0,1\}^{\kappa}$; $\boldsymbol{R} = f_R(\rho)$	6. $\boldsymbol{R} \xleftarrow{\$} \chi_S^{\overline{m}}$
7. $\boldsymbol{U} = \mathrm{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$	7. $\boldsymbol{U} = \mathrm{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$
8. $\boldsymbol{V} = \langle (\mathrm{R}_{p \to t, h_2}(\langle \boldsymbol{B}^T \boldsymbol{R} \rangle_{\xi})) + \frac{t}{b} \boldsymbol{M_b} \rangle_t$	8. $\boldsymbol{V} = \langle (\mathrm{R}_{p \to t, h_2}(\langle \boldsymbol{B}^T \boldsymbol{R} \rangle_{\xi})) + \frac{t}{b} \boldsymbol{M_b} \rangle_t$
9. $\beta' = \mathcal{A}\left((\boldsymbol{A}, \boldsymbol{B}), (\boldsymbol{U}, \mathrm{Sample}_\mu(\boldsymbol{V})), st\right)$	9. $\beta' = \mathcal{A}\left((\boldsymbol{A}, \boldsymbol{B}), (\boldsymbol{U}, \mathrm{Sample}_\mu(\boldsymbol{V})), st\right)$
10. $\mathbf{return}[(\beta' = \beta)]$.	10. $\mathbf{return}[(\beta' = \beta)]$.

As a result, the pairs $(\boldsymbol{B}, \boldsymbol{v})$ in Game G_5 and $(\langle \boldsymbol{B_q} \rangle_p, \langle \boldsymbol{v'} \rangle_t)$ in Game G_6 have the same distribution, and so

$$\Pr(S_5) = \Pr(S_6). \tag{44}$$

We now consider games G_6 and G_7. We will use the following lemma.

Lemma A.4.0.1. *Let a, b, c be positive integers such that $a|b|c$. Then for any $x \in \mathbb{R}$*

$$\lfloor \frac{a}{c} x \rceil = \lfloor \frac{a}{b} \lfloor \frac{b}{c} x \rceil \rceil$$

Proof Write $m = \frac{c}{b}$ and $n = \frac{b}{a}$. It is sufficient that to show that

$$\lfloor \frac{x}{mn} \rceil = \lfloor \frac{1}{n} \lfloor \frac{x}{m} \rceil \rceil.$$

We write $x = mn\lfloor \frac{x}{mn} \rceil + y$ with $0 \le y < mn$. Obviously, $\frac{1}{n} \lfloor \frac{x}{m} \rceil = \lfloor \frac{x}{mn} \rceil + \frac{1}{n} \lfloor \frac{y}{m} \rceil$. As $0 \le y < mn$, it holds that $0 \le \frac{y}{m} < n$ and so $0 \le \lfloor \frac{y}{m} \rceil \le n - 1$. $\square$

Applying the above lemma, we infer that $(\boldsymbol{U}, \boldsymbol{V'})$ in Game G_6 and $(\boldsymbol{U'}, \boldsymbol{V"})$ in Game G_7 are related as $\boldsymbol{U} = \lfloor \frac{p}{z} \boldsymbol{U'} \rceil$ and $\boldsymbol{V'} \equiv \lfloor \frac{tq}{pz} \boldsymbol{V"} \rceil \pmod{\frac{tq}{p}}$. As a result,

$$\Pr(S_6) = \Pr(S_7). \tag{45}$$

In Game G_8, the variable $\begin{bmatrix} \mathrm{R}_{q \to z, h_2}(\langle \boldsymbol{A^T R} \rangle_{\phi_{n+1}}) \\ \mathrm{R}_{q \to z, h_2}(\langle \boldsymbol{B^T R} \rangle_{\xi}) \end{bmatrix}$ from Game G_7 is replaced by the $(d/n + \overline{n}) \times \overline{m}$ matrix $\begin{bmatrix} \boldsymbol{U"} \\ \boldsymbol{W'} \end{bmatrix}$ with entries uniformly drawn from $\mathcal{R}_{n,z}$.

Table 25: IND-CPA games for r5_cpa_pke: games G_6, G_7 and G_8

Game G_6	Game G_7	Game G_8
1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$	1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$	1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$
2. -	2. -	
3. $\boldsymbol{B_q} \xleftarrow{\$} \mathcal{R}_{n,p}^{d/n \times \overline{n}}$	3. $\boldsymbol{B_q} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times \overline{n}}$	3. $\boldsymbol{B_q} \xleftarrow{\$} \mathcal{R}_{n,p}^{d/n \times \overline{n}}$
4. Choose $\beta \xleftarrow{\$} \{0,1\}$.	4. Choose $\beta \xleftarrow{\$} \{0,1\}$.	4. Choose $\beta \xleftarrow{\$} \{0,1\}$.
5. $(m_0, m_1, st) = \mathcal{A}(\boldsymbol{A}, \langle \boldsymbol{B_q} \rangle_p)$.	5. $(m_0, m_1, st) = \mathcal{A}(\boldsymbol{A}, \langle (\boldsymbol{B_q}) \rangle_p)$.	5. $(m_0, m_1, st) = \mathcal{A}(\boldsymbol{A}, \langle \boldsymbol{B_q} \rangle_p)$.
6. $\boldsymbol{R} \xleftarrow{\$} \chi_S^{\overline{m}}$	6. $\boldsymbol{R} \xleftarrow{\$} \chi_S^{\overline{m}}$	
7. $\boldsymbol{U} = \text{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$	7. $\boldsymbol{U'} = \text{R}_{q \to z, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$ $\quad z = \max(p, tq/p)$	7. $\boldsymbol{U"} \xleftarrow{\$} \mathcal{R}_{n,z}^{d/n \times \overline{m}}$
8. $\boldsymbol{V'} = \langle (\text{R}_{q \to tq/p, h_2}(\langle \boldsymbol{B_q}^T \boldsymbol{R} \rangle_\xi)) + \frac{t}{b} \boldsymbol{M_b} \rangle_{tq/p}$	8. $\boldsymbol{V''} = \langle (\text{R}_{q \to z, h_2}(\langle \boldsymbol{B_q}^T \boldsymbol{R} \rangle_\xi)) + \frac{pz}{qb} \boldsymbol{M_b} \rangle_z$	8. $\boldsymbol{W} \xleftarrow{\$} \mathcal{R}_{n,z}^{\overline{n} \times \overline{m}}$; $\boldsymbol{V'''} = \langle \boldsymbol{W} + \frac{pz}{qb} \boldsymbol{M_b} \rangle_z$
9. $\beta' = \mathcal{A}\left((\boldsymbol{A}, \langle \boldsymbol{B_q} \rangle_p), (\boldsymbol{U}, \text{Sample}_\mu(\langle \boldsymbol{V'} \rangle_t)), st\right)$	9. $\beta' = \mathcal{A}((\boldsymbol{A}, \langle \boldsymbol{B_q} \rangle_p), (\langle \lfloor \frac{p}{z} \boldsymbol{U'} \rceil \rangle_p,$ $\text{Sample}_\mu(\langle \lfloor \frac{tq}{pz} \boldsymbol{V''} \rceil \rangle_t), st)$	9. $\beta' = \mathcal{A}((\boldsymbol{A}, \langle \boldsymbol{B_q} \rangle_p), (\langle \lfloor \frac{p}{z} \boldsymbol{U"} \rceil \rangle_p,$ $\text{Sample}_\mu(\langle \lfloor \frac{tq}{pz} \boldsymbol{V'''} \rceil \rangle_t), st)$
10. $\textbf{return}[(\beta' = \beta)]$.	10. $\textbf{return}[(\beta' = \beta)]$.	10. $\textbf{return}[(\beta' = \beta)]$.

As $h_2 = \frac{q}{2z}$, $\text{R}_{q \to z, h_2} = \text{R}_{q \to z}$. Moreover, if $\xi = \Phi_{n+1}$, the first variable in fact equals $\text{R}_{q \to z}(\langle \begin{bmatrix} \boldsymbol{A}^T \\ \boldsymbol{B}^T \end{bmatrix} \boldsymbol{R} \rangle_\phi)$. We infer that there exists a distinghuisher $\mathcal{E}$ such that

$$\text{Adv}_{d,n,d/n+\overline{n},q,z}^{\text{dGLWR}_{\text{spt}}}(\mathcal{A} \circ \mathcal{Z}) = |\Pr(S_7) - \Pr(S_8)| \tag{46}$$

As all inputs to $\mathcal{A}$ in Game G_8 are uniform, $\Pr(S_8) = \frac{1}{2}$, and so

$$\text{Adv}_{\text{CPA-PKE}}^{\text{IND-CPA}}(\mathcal{A}) = |\Pr(S_0) - \Pr(S_8)| = |\sum_{i=0}^{7} \Pr(S_i) - \Pr(S_{i+1})| \leq \sum_{i=0}^{7} |\Pr(S_i) - \Pr(S_{i+1})|. \tag{47}$$

By combining (38)-(47), the theorem follows.

A.5   IND-CPA Security for RLWE-based Round5 variant with different reduction polynomials

As described in Section 1.4.1, the merged parameters of Round5 require that the reduction polynomial $\xi(x)$ used in the computation of the ciphertext component $\boldsymbol{v}$ and in decryption equals $x^{n+1} - 1$. The proof of the IND-CPA security of r5_cpa_pke presented in Section A.4, however, only applies if $\xi(x) = \Phi_{n+1}(x)$. The reason is that otherwise the replacement of $(\boldsymbol{U'}, \boldsymbol{V''})$ by random matrices in the transition from game G_7 to game G_8 cannot be directly related to the difficulty of GLWR with one single reduction polynomial.

Next we argue why this construction is secure: First, Round5 uses function $Sample_\mu$ that selects μ coefficients out of $n + 1$. We show how it prevents

known distinguishing attacks such as the "*Evaluate at 1*" attack [55]. Second, we discuss an extension of the IND-CPA security proof in Section A.4 for a RLWE-variant of Round5. The existence of this proof for the RLWE case gives strong confidence in the Round5 parameters using error correction. Finally, we discuss why this proof does not directly translate to an RLWR based design and a simple design change in Round5 that would make it work, but that we have not introduced since it does not bring major benefits from a concrete security viewpoint.

Distinguishing attack at $x = 1$. In this section, the well-known "Evaluate at $x = 1$" distinguishing attack [55] is described that can be applied if $\xi(x) = x^{n+1} - 1$ and $\mu = n + 1$. Next, it is argued that this attack cannot be applied in Round5 if $\mu \leq n$. As a shorthand, $N(x)$ is written instead of $x^{n+1} - 1$.

Consider a pair of polynomials $(b(x), v(x))$ with $b(x)$ uniformly distributed on $\mathbb{Z}_p[x]/(x^{n+1}-1)$ and $v(x) = \langle \text{Sample}_\mu(\lfloor \frac{t}{p}(\langle b(x)r(x)\rangle_{N(x)} + h_2)\rceil) + \frac{t}{b}m(x)\rangle_t$ with $r(x)$ drawn independently and uniformly from the ternary polynomials of degree at most $n-1$ satisfying $r(1) = 0$, and $m(x)$ drawn according to some distribution on $\mathbb{Z}_b[x]/(x^\mu - 1)$. We then have that $v(x) \equiv \lfloor \text{Sample}_\mu(\frac{t}{p}(\langle b(x)r(x)\rangle_{N(x)} + h_2)\rceil) + \frac{t}{b}m(x) \pmod{t}$, and so $w(x) = \frac{p}{t}v(x)$ satisfies

$$w(x) \equiv \text{Sample}_\mu(\langle b(x)r(x)\rangle_{N(x)}) + \frac{p}{t} \cdot h_2 \sum_{i=0}^{\mu-1} x^i - \frac{p}{t}\epsilon(x) + \frac{p}{b}m(x) \pmod{p}.$$

where $\epsilon(x)$ is the result of rounding downwards, so all components of $\frac{p}{t}\epsilon(x)$ are in $[0, \frac{p}{t}) \cap \mathbb{Z}$. As $(x-1)$ divides both $r(x)$ and $N(x)$, it follows that $x-1$ divides $\langle b(x)r(x)\rangle_{N(x)}$, and so if $\mu = n + 1$, then

$$w(1) \equiv \frac{p}{t} \cdot h_2 \cdot (n+1) - \frac{p}{t}\sum_{i=0}^{n}\epsilon_i + \frac{p}{b}m(1) \pmod{p}.$$

For large n, the value of $\frac{p}{t}\sum_{i=0}^{n}\epsilon_i$ is close to its average, i.e., close to $n\frac{p}{2t}$. As a result, has maxima at values $\frac{p}{t}h_2(n+1) - n\frac{p}{2t} + \frac{p}{b}k$ for $0 \leq k \leq b-1$. So $w(1)$ can serve as a distinguisher between the above distribution and the uniform one. Now assume that $\mu < n + 1$. We take $\mu = n$, which is the case giving most information to the attacker. Writing $f(x) = \langle b(x)r(x)\rangle_{N(x)} = \sum_{i=0}^{n}f_ix^i$, it holds that

$$w(1) \equiv \sum_{i=0}^{n-1}f_i + \frac{p}{t} \cdot h_2 \cdot n - \frac{p}{t}\epsilon(1) + \frac{p}{b}m(1) \pmod{p}.$$

As shown above, $f(1) = 0$, and so $\sum_{i=0}^{n-1}f_i = -f_n$. Hence, under the assumption that f_n is distributed uniformly modulo p, also $w(1)$ is distributed uniformly modulo p. The latter assumption is supported by [79].

Requirements for IND-CPA Security – Design rationale. An IND-CPA Security proof is feasible for a RLWE variant of r5_cpa_pke, i.e., a Round5 variant where the noise is independently generated. This proof is presented below and gives confidence in Round5's design and choices made.

The proof requires the secrets to be a multiple of $(x - 1)$ and also the noise polynomial for the ciphertext component v to be a multiple of $(x - 1)$ (this is used in an essential step of the proof, specifically in the map Ψ in (56)) This last requirement is the reason why this proof does not apply to Round5 with $\xi(x) = x^{n+1} - 1$ using RLWR defined as *component-wise rounding*. This deterministic component-wise rounding does not allow enforcing that the noisy "rounding" polynomials are multiples of $(x - 1)$.

Round5's design can be adapted to use a slightly different type of rounding informally named as "rounding to the root lattice"[1] - that allows the IND-CPA proof to work. This "rounding to the root lattice" would work in three steps: (1) perform component-wise rounding as currently done by rounding (down) from a higher modulus q to a smaller one p; (2) compute the difference k between the sum of the rounded coefficients and the next multiple of p; and (3) add 1 to the k rounded smallest coefficients. The components of the rounding noise computed in this way adds up to "0" so that the IND-CPA proof works.

However, this modification – going from component-wise rounding to rounding to the root lattice – would introduce additional complexity and changes in Round5's design with no clear concrete security benefits. First, Sample$_\mu$ gets rid of $n + 1 - \mu$ coefficients so that knowing k is irrelevant. Second, concrete security attacks use the norm of the noise that hardly changes if rounding to the root lattice is employed. Because of these two reasons, we argue that the current Round5 design (and the rounding used in it) is sound and secure, and further modifications are not required.

IND-CPA Proof for RLWE-based variant of Round5. We now present the proof of IND-CPA security for the RLWE variant of r5_cpa_pke. We restrict ourselves to the case $B = 1$, that is, one single bit is extracted from each symbol.

The following notation will be used. We write $\phi(x) = 1 + x + \ldots + x^n$, and $N(x) = x^{n+1} - 1$, where $n + 1$ is prime. Moreover, $R_\phi = \mathbb{Z}_q[x]/\phi(x)$, and

$$R_0 = \{f(x) = \sum_{i=0}^{n} f_i x^i \in \mathbb{Z}_q[x] \mid \sum_{i=0}^{n} f_i \equiv 0 \pmod{q}\} \tag{48}$$

As $N(x) = (x - 1)\phi(x)$, it holds that $\langle (x-1)f(x) \rangle_{N(x)} = (x-1)\langle f(x) \rangle_{\phi(x)}$ for any $f \in \mathbb{Z}[x]$. As a result, $f(x) \mapsto (x-1)f(x)$ is a bijection from R_ϕ to R_0.

In the proof, the following lemma will be used.

Lemma A.5.0.1. *Let q and $n + 1$ be relatively prime, and let $(n + 1)^{-1}$ be the*

---

[1] We thank Léo Ducas for helpful discussions on this topic.

multiplicative inverse of $n+1$ in $\mathbb{Z}_q$. The mapping $\mathcal{F}$ defined as

$$\mathcal{F} : (\sum_{i=0}^{n-1} f_i x^i) \mapsto \sum_{i=0}^{n-1} f_i x^i - (n+1)^{-1} \cdot \left(\sum_{i=0}^{n-1} f_i\right) \cdot \phi(x)$$

is a bijection from R_ϕ to R_0.

Proof. It is easy to see that $\mathcal{F}$ maps R_ϕ to R_0. To show that $\mathcal{F}$ is a bijection, let $g(x) = \sum_{i=0}^{n} g_i x^i \in R_0$, and let $f(x) = \sum_{i=0}^{n} \langle g_i - g_n \rangle_q x^i$. Clearly, $f \in \mathbb{Z}_q[x]$ has degree at most $n-1$, and by direct computation, $\mathcal{F}(f(x)) = g(x)$. $\qquad\square$

In the description below, $\mathcal{S}$ denotes a set of secrets such that

$$\mathcal{S} \subset \{f(x) = \sum_{i=0}^{n-1} f_i x^i \in \mathbb{Z}_q[x] \mid \sum_{i=0}^{n-1} f_i \equiv 0 \pmod{q}\}, \qquad (49)$$

Moreover, $\mathcal{M}$ denotes a message space, and ECC_Enc and ECC_Dec are error correcting encoding and decoding algorithms such that

$$\{ECC\_Enc(m) \mid m \in \mathcal{M}\} \subset \{f(x) = \sum_{i=0}^{n} f_i x^i \in \mathbb{Z}_2[x] \mid \sum_{i=0}^{n} f_i \equiv 0 \pmod{2}\}.$$
$$(50)$$

Moreover, χ denotes a probability distribution on R_ϕ.

For understanding Algorithm 24 below, note that as $(x-1)|s(x)$, we have that $su' \equiv sa'r + se_1 \pmod{N}$, and, as $(x-1)|r(x)$, that $rb' \equiv ra's + re_0 \pmod{N}$. As a consequence,

$$\zeta \equiv v - su' \equiv \frac{q}{2} ECC\_Enc(m) + (x-1)e_2 + re_0 - se_1 \pmod{N}, \text{ whence}$$

$$\lfloor \frac{2}{q}\zeta \rceil \equiv ECC\_Enc(m) + \lfloor \frac{2}{q}((x-1)e_2 + re_0 - se_1) \rceil \pmod{N}.$$

Algorithm 22: CPA-PKE.Keygen()

---

1 $a' \xleftarrow{\$} R_\phi, s \xleftarrow{\$} \mathcal{S}, e_0 \leftarrow \chi$
2 $b' = \langle a's + e_0 \rangle_\phi$
3 $pk = (a', b')$
4 $sk = s$
5 **return** (pk, sk)

We are now in a position to prove the following result.

Theorem A.5.0.1. *For every IND-CPA adversary $\mathcal{A}$ with advantage A, there exist algorithms C and E such that*

$$A \leq Adv_1(C) + Adv_3(E). \qquad (51)$$

Algorithm 23: CPA-PKE.Enc$(pk = (a', b'), m \in \mathcal{M})$

---

1 $r \xleftarrow{\$} \mathcal{S}, e_1, e_2 \xleftarrow{\$} \chi$

2 $u' = \langle a'r + e_1 \rangle_\phi$

3 $v = \langle \frac{q}{2} ECC\_Enc(m) + b'r + (x-1)e_2 \rangle_N$

4 $ct = (u', v)$

5 return ct

Algorithm 24: CPA-PKE.Dec(sk, ct)

---

1 $\zeta = \langle v - su' \rangle_N$

2 $\hat{m} = ECC\_Dec\langle \lfloor \frac{2\zeta}{q} \rceil \rangle_2)$

3 return $\hat{m}$

Here Adv_1 refers to the advantage of distinguishing between the uniform distribution on $(\mathbb{Z}_q[x]/\phi(x))^2$ and the R-LWE distribution

$$(a', b' = \langle a's + e_0 \rangle_\phi) \text{ with } a' \xleftarrow{\$} R_\phi, s \xleftarrow{\$} \mathcal{S}, e_0 \leftarrow \chi \qquad (52)$$

Similarly, Adv_3 refers to the advantage of distinguishing between the uniform distribution on $(\mathbb{Z}_q[x]/\phi(x))^4$ and the distribution of two R-LWE samples with a common secret, given by

$$(a', b'', u', v') \text{ with } a', b'' \xleftarrow{\$} \mathbb{Z}_q[x]/\phi(x), u = \langle a'r + e_1 \rangle_\phi, \qquad (53)$$

$$v = \langle b''r + e_2 \rangle_\phi \text{ with } r \xleftarrow{\$} \mathcal{S}, e_1, e_2 \leftarrow \chi \qquad (54)$$

Proof. We prove the theorem using a sequence of IND-CPA games. We denote by S_i the event that the output of game i equals 1.

Game G_0 is the original IND-CPA game. In Game G_1, the public key (a', b') is replaced by a pair (a', b') uniformly drawn from R_ϕ^2. It can be shown that there exists an algorithm $\mathcal{C}$ for distinguishing between the uniform distribution on R_ϕ^2 and the R-LWE distribution of pairs (a', b') with $a' \xleftarrow{\$} R_\phi$, $b' = \langle as' + e_0 \rangle_\phi$ with $s \xleftarrow{\$} \mathcal{S}$ and $e_0 \leftarrow \chi$ such that

$$\text{Adv}_1(\mathcal{C}) = |\text{Pr}(S_0) - \text{Pr}(S_1)|.$$

In Game G_2, the values $u' = \langle a'r + e_1 \rangle_\phi$ and $\hat{v} = \langle b'r + (x-1)e_2 \rangle_N$ used in the generation of v are simultaneously substituted with uniform random variables from R_ϕ and R_0, respectively. it can be shown that there exists an adversary $\mathcal{D}$ with the same running time as that of $\mathcal{A}$ such that

$$\text{Adv}_2(\mathcal{D}) = |\text{Pr}(S_1) - \text{Pr}(S_2)|.$$

Here Adv_2 refers to the advantage of distinguishing between the uniform distribution on $R_\phi^3 \times R_0$ and the distribution

$$(a', b', u', v) = (a', b', \langle a'r + e_1 \rangle_\phi, \langle b'r + (x-1)e_2 \rangle_N) \text{ with } a', b' \xleftarrow{\$} R_\phi, r \xleftarrow{\$} \mathcal{S}, e_1, e_2 \xleftarrow{\$} \chi.$$
(55)

Because of (50), the value of the ciphertext v in Game G_2 is independent of bit b, and therefore $\Pr(S_2) = 1/2$. As a final step, we define $\Psi : R_\phi^3 \times R_0 \to R_\phi^4$ as

$$\Psi(a'(x), b'(x), u'(x), v(x)) = (a'(x), b''(x), u'(x), v'(x)) \text{ with} \qquad (56)$$

$$b''(x) = \frac{\mathcal{F}(b'(x))}{x-1}, v'(x) = \frac{v(x)}{x-1} \qquad (57)$$

As $\mathcal{F}$ is a bijection from R_ϕ to R_0 (see Lemma A.5.0.1) and $f(x) \mapsto \frac{f(x)}{x-1}$ is a bijection from R_0 to R_ϕ, it follows that Ψ is a bijection. Writing $b(x) = \mathcal{F}(b'(x))$, we infer that

$$b(x)r(x) = b'(x)r(x) - (n+1)^{-1}b'(1)\phi(x)r(x) \equiv b'(x)r(x) \pmod{N(x)},$$

where the latter equivalence holds as $r(x)$ is a multiple of $(x-1)$, and so

$$v(x) = \langle b'(x)r(x) + (x-1)e_2(x) \rangle_N = \langle b(x)r(x) + (x-1)e_2(x) \rangle_N.$$

As $r(x)$ is a multiple of $x-1$, it follows that $v(x) \in R_0$ and that

$$v'(x) = \frac{v(x)}{x-1} \equiv \langle b''(x)r(x) + e_2(x) \rangle_\phi \text{ where } b''(x) = \frac{b(x)}{x-1}.$$

As a result, the advantage of $\mathcal{E} = \Psi \circ \mathcal{D}$ in distinguishing between the uniform distribution on R_Φ^4 and the distribution

$$(a', b'', u', v') \text{ with } a, b'' \xleftarrow{\$} R_\phi, u'(x) = \langle a'r + e_1 \rangle_\phi \text{ and } v' = \langle b''r + e_2 \rangle_\phi$$

is equal to $\text{Adv}_2(D)$. Note that (a, u') and $(b", v')$ are two R-LWE samples with common secret $r(x) \in \mathcal{S}$, with $a', b"$ chosen uniformly in $\mathcal{R}_\phi$ and independent noise polynomials $e_1(x)$ and $e_2(x)$.
As $Pr(S_2) = \frac{1}{2}$, we conclude that

$$Adv(\mathcal{A}) = |\Pr(S_0) - \Pr(S_2)| \leq \sum_{i=0}^{1} |\Pr(S_i) - \Pr(S_{i+1})| = \text{Adv}_1(\mathcal{C}) + \text{Adv}_2(\mathcal{E}).$$

$$\square$$

A.6   IND-CPA Security of r5_cpa_kem

This section presents a proof that r5_cpa_kem is IND-CPA secure, based on the hardness of the decision GLWR problem with sparse-ternary secrets. In the proof, instead of using $pk = (\sigma, \boldsymbol{B})$, we use $pk = (\boldsymbol{A} = f_{d,n}^\tau(\sigma), \boldsymbol{B})$. The distribution of $f_{d,n}^\tau(\sigma)$ with $\sigma \xleftarrow{\$} \{0,1\}^\kappa$ is denoted by F_τ. Moreover, the uniform distribution $\mathcal{H}_{n,d/n}(h)$ is denoted by χ_S. The proof is restricted to the case $\xi(x) = \Phi_{n+1}(x)$.

Theorem A.6.0.1. *Let b, p, q, t be powers of two such that $b|t|p|q$, and let $z = \max(p, tq/p)$. Furthermore, assume that $\xi(x) = \Phi_{n+1}(x)$. If $f_{d,n}^{(\tau)}$, f_S and f_R induce distributions indistinguishable from uniform, and H is a secure pseudo-random function, then r5_cpa_kem is IND-CPA secure under the hardness assumption of the Decision-GLWR problem with sparse-ternary secrets. More precisely, if $Adv_{r5\_cpa\_kem}^{IND\text{-}CPA}(\mathcal{A})$ is the advantage of adversary $\mathcal{A}$ in distinguishing a key encapsulated using r5_cpa_kem from random, then there exist distinghuishers $\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}, \mathcal{F}, \mathcal{G}$ such that*

$$
\begin{aligned}
Adv_{r5\_cpa\_kem}^{IND\text{-}CPA}(\mathcal{A}) &\leq Adv_{\mathcal{U}(\mathcal{R}_{n,q}^{d/n \times d/n}), F_\tau}(\mathcal{B}) + Adv_{G_S, \chi_S^{\overline{n}}}(\mathcal{C}) + Adv_{G_R, \chi_S^{\overline{m}}}(\mathcal{D}) \\
&+ \overline{n} \cdot Adv_{d,n,d/n,q,p}^{dGLWR_{spt}}(\mathcal{E}) + Adv_{d,n,d/n+\overline{n},q,z}^{dGLWR_{spt}}(\mathcal{F}) + Adv_{G_H, \mathcal{U}(\{0,1\}^\kappa)}(\mathcal{G})
\end{aligned}
\tag{58}
$$

In this equation, F_τ is the distribution of $f_{d,n}^{(\tau)}$ with $\sigma \overset{\$}{\leftarrow} \{0,1\}^\kappa$, G_S is the distribution of $f_S(s)$ with $s \overset{\$}{\leftarrow} \{0,1\}^\kappa$ and G_R is the distribution of $f_R(\rho)$ with $\rho \overset{\$}{\leftarrow} \{0,1\}^\kappa$. Moreover, $Adv_{d,n,m,q_1,q_2}^{dGLWR_{spt}}(\mathcal{Z})$ is the advantage of adversary $\mathcal{Z}$ in distinguishing m GLWR samples (with sparse-ternary secrets) from uniform, with the GLWR problem defined for the parameters d, n, q_1, q_2. Finally, G_H is the distribution of $H(x)$ with $x \overset{\$}{\leftarrow} \{0,1\}^{\kappa+\lambda_1}$ with $\lambda_1 = d \cdot \overline{n} \log_2(p) + \mu \log_2(t)$.

Proof. The proof for Theorem A.6.0.1, proceeds via a similar sequence of games as in the proof of Theorem A.4.0.1, shown in Tables 26 to 29. Again, S_i denotes the event that the output in game G_i equals 1. As the sequence of games is essentially the same as for the proof of Theorem A.4.0.1, we focus on the essential difference, which is game G_8. Game G_8 differs from game G_7 only in the generation of K_0. As p, q, t all are powers of two and $z = \max(\frac{tq}{p}, p)$, p divides z and tq divides pz. As a result, in games G_7 and G_8, $\lfloor \frac{p}{z} U'' \rceil$ is uniformly distributed on $\mathcal{R}_{n,p}^{d/n \times \overline{m}}$. Also, as W''' in games G_7 and G_8 is uniformly distributed on $\mathcal{R}_{n,z}^{\overline{n} \times \overline{m}}$, the vector $\lfloor \frac{tq}{pz} v''' \rceil$ is uniformly distributed on $\mathbb{Z}_{\frac{tq}{p}}$. As t divides $\frac{tq}{p}$, the vector $\langle \lfloor \frac{tq}{pz} v''' \rceil \rangle_t$ is uniformly distributed on $\mathbb{Z}_t$. We conclude that the input to H in game G_7 is uniform. Therefore, there exists a distinguisher between the uniform distribution on $\{0,1\}^\kappa$ and the distribution of $H(x)$ with $x \overset{\$}{\leftarrow} \{0,1\}^{\kappa+\lambda_1}$ (where $\lambda_1 = d \cdot \overline{n} \log_2(p) + \mu \log_2(t)$) with advantage equal to $|\Pr(S_7) - \Pr(S_8)|$. As the input to $\mathcal{A}$ in game G_8 is uniform, $\Pr(S_8) = \frac{1}{2}$. $\square$

A.7   IND-CCA security of r5_cca_pke

In this section, it is shown that r5_cca_pke is IND-CCA secure. As r5_cca_pke is constructed from r5_cca_kem and a secure data-encapsulation mechanism as proposed by Cramer and Shoup [39], it is sufficient to show the IND-CCA security of r5_cca_kem. Indeed, as stated in Theorem A.7.0.1, when the hash functions G and H in Algorithms 8 and 9 are modeled as random oracles, the key-encapsulation mechanism r5_cca_kem defined in Section 1.4.4 is IND-CCA secure, assuming the hardness of the decision GLWR problem with sparse-ternary secrets.

Table 26: IND-CPA games for r5_cpa_kem: Games G_0 and G_1

Game G_0	Game G_1
1. $\boldsymbol{A} \leftarrow F_\tau$	1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$
2. $s \xleftarrow{\$} \{0,1\}^\kappa$; $\boldsymbol{S} = f_S(s)$	2. $s \xleftarrow{\$} \{0,1\}^\kappa$; $\boldsymbol{S} = f_S(s)$;
3. $\boldsymbol{B} = \mathrm{R}_{q \to p}(\langle \boldsymbol{AS} \rangle_{\Phi_{n+1}})$	3. $\boldsymbol{B} = \mathrm{R}_{q \to p}(\langle \boldsymbol{AS} \rangle_{\Phi_{n+1}})$
4. Choose $b \xleftarrow{\$} \{0,1\}$.	4. Choose $b \xleftarrow{\$} \{0,1\}$.
5. $(ct = (\boldsymbol{U}, \boldsymbol{v}), K_0) = \mathrm{r5\_cpa\_kem\_encapsulate}\,(\boldsymbol{A}, \boldsymbol{B})$.	5. $(ct = (\boldsymbol{U}, \boldsymbol{v}), K_0) = \mathrm{r5\_cpa\_kem\_encapsulate}\,(\boldsymbol{A}, \boldsymbol{B})$.
6. $K_1 \xleftarrow{\$} \{0,1\}^\kappa$.	6. $K_1 \xleftarrow{\$} \{0,1\}^\kappa$.
7. $b' = \mathcal{A}\,((\boldsymbol{A}, \boldsymbol{B}), ct, K_b)$.	7. $b' = \mathcal{A}\,((\boldsymbol{A}, \boldsymbol{B}), ct, K_b)$.
8. Output $[(b' = b)]$.	8. Output $[(b' = b)]$.

Table 27: IND-CPA games for r5_cpa_kem: Games G_2 and G_3

Game G_2	Game G_3
1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$	1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$
2. $\boldsymbol{S} \leftarrow \chi_S^{\overline{n}}$	2. -
3. $\boldsymbol{B} = \mathrm{R}_{q \to p}(\langle \boldsymbol{AS} \rangle_{\Phi_{n+1}})$	3. $\boldsymbol{B} \xleftarrow{\$} \mathcal{R}_{n,p}^{d/n \times \overline{n}}$
4. Choose $b \xleftarrow{\$} \{0,1\}$.	4. Choose $b \xleftarrow{\$} \{0,1\}$.
5. $(ct = (\boldsymbol{U}, \boldsymbol{v}), K_0) = \mathrm{r5\_cpa\_kem\_encapsulate}\,(\boldsymbol{A}, \boldsymbol{B})$.	5. $(ct = (\boldsymbol{U}, \boldsymbol{v}), K_0) = \mathrm{r5\_cpa\_kem\_encapsulate}\,(\boldsymbol{A}, \boldsymbol{B})$.
6. $K_1 \xleftarrow{\$} \{0,1\}^\kappa$.	6. $K_1 \xleftarrow{\$} \{0,1\}^\kappa$.
7. $b' = \mathcal{A}\,((\boldsymbol{A}, \boldsymbol{B}), ct, K_b)$.	7. $b' = \mathcal{A}\,((\boldsymbol{A}, \boldsymbol{B}), ct, K_b)$.
8. Output $[(b' = b)]$.	8. Output $[(b' = b)]$.

Theorem A.7.0.1. *For any adversary $\mathcal{A}$ that makes at most q_H queries to the random oracle H, at most q_G queries to the random oracle G, and at most q_D queries to the decryption oracle, there exists an adversary $\mathcal{B}$ such that*

$$Adv_{CCA\text{-}KEM}^{IND\text{-}CCA}(\mathcal{A}) \leq 3 \cdot Adv_{CPA\text{-}PKE}^{IND\text{-}CPA}(\mathcal{B}) + q_G \cdot \delta + \frac{2q_G + q_H + 1}{2^{\mu B}} \quad (59)$$

when r5_cpa_pke and r5_cca_kem both have a probability of decryption/decapsulation failure that is at most δ.

Proof. The proof of Theorem A.7.0.1 proceeds via two transformation reductions due to [53]. First, Lemma A.7.0.1 establishes that the OW-PCA [2] security of the deterministic public-key encryption scheme PKE₁ obtained from the public-key encryption scheme PKE via transformation T [53], tightly reduces to IND-CPA security of PKE₁. This lemma is a special case of [53, Theorem 3.2]

---
[2]The security notion of <u>O</u>ne-<u>W</u>ay against <u>P</u>laintext <u>C</u>hecking <u>A</u>ttacks.

Table 28: IND-CPA games for r5_cpa_kem: Games G_4 and G_5

Game G_4	Game G_5
1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$, $\boldsymbol{B} \xleftarrow{\$} \mathcal{R}_{n,p}^{d/n \times \overline{n}}$.	1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$, $\boldsymbol{B}_q \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times \overline{n}}$.
2. Choose $b \xleftarrow{\$} \{0,1\}$.	2. Choose $b \xleftarrow{\$} \{0,1\}$.
3. $m \xleftarrow{\$} \{0,1\}^{\kappa}, \boldsymbol{\zeta} = ECC\_Enc_{\kappa,f}(m)$	3. $m \xleftarrow{\$} \{0,1\}^{\kappa}, \boldsymbol{\zeta} = ECC\_Enc_{\kappa,f}(m)$
4. $\boldsymbol{R} \xleftarrow{\$} \left(\mathcal{H}_{n,d/n}(h)\right)^{1 \times \overline{m}}$.	4. $\boldsymbol{R} \xleftarrow{\$} \left(\mathcal{H}_{n,d/n}(h)\right)^{1 \times \overline{m}}$
5. $\boldsymbol{U} = \mathrm{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$	5. $\boldsymbol{U} = \mathrm{R}_{q \to p, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$
6. $\boldsymbol{W} = \mathrm{R}_{p \to t, h_2}(\langle \boldsymbol{B}^T \boldsymbol{R} \rangle_{\xi})$	6. $\boldsymbol{W}' = \mathrm{R}_{q \to tq/p, h_2}(\langle \boldsymbol{B}_q^T \boldsymbol{R} \rangle_{\xi})$
7. $\boldsymbol{v} = \langle \mathrm{Sample}_{\mu}(\boldsymbol{W}) + \frac{t}{b}\boldsymbol{\zeta} \rangle_t$	7. $\boldsymbol{v}' = \langle \mathrm{Sample}_{\mu}(\boldsymbol{W}') + \frac{t}{b}\boldsymbol{\zeta} \rangle_{tq/p}$
8. $K_0 = H(m, \mathtt{bin}_1(ct = (\boldsymbol{U}, \boldsymbol{v})))$.	8. $K_0 = H(m, \mathtt{bin}_1(ct = (\boldsymbol{U}, \langle \boldsymbol{v}' \rangle_t)))$.
9. $K_1 \xleftarrow{\$} \{0,1\}^{\kappa}$.	9. $K_1 \xleftarrow{\$} \{0,1\}^{\kappa}$.
10. $b' = \mathcal{A}((\boldsymbol{A}, \boldsymbol{B}), ct, K_b)$.	10. $b' = \mathcal{A}((\boldsymbol{A}, \langle \boldsymbol{B}_q \rangle_p), ct, K_b)$.
11. Output $[(b' = b)]$.	11. Output $[(b' = b)]$.

with $q_v = 0$, since by definition OW-PCA security is OW-PCVA[3] security where the attacker is not allowed to query the ciphertext validity checking oracle.

Lemma A.7.0.1 (Adapted from [53, Theorem 3.2]). *Assume PKE to be δ correct. Then, for any OW-PCA adversary $\mathcal{B}$ that issues at most q_G queries to the random oracle G, q_P queries to a plaintext checking oracle P_{CO}, there exists an IND-CPA adversary $\mathcal{C}$ such that*

$$Adv_{PKE_1}^{OW\text{-}PCA}(\mathcal{B}) \leq q_G \cdot \delta + \frac{2q_G + 1}{|\mathcal{M}|} + 3 \cdot Adv_{PKE}^{IND\text{-}CPA}(\mathcal{C}) \tag{60}$$

where $\mathcal{M}$ is the message/plaintext space of the public-key encryption schemes PKE and PKE_1.

Next, combinination of Lemma A.7.0.1 and the reduction in [53, Theorem 3.4] shows that the IND-CCA security of a KEM with implicit rejection that is constructed using a non-deterministic PKE (like r5_cca_kem), tightly reduces to the IND-CPA security of said PKE. □

Direct application of [53, Theorem 4.6], similarly as in [26, Theorem 4.2], shows that r5_cca_kem is IND-CCA secure in the quantum random oracle model. The resulting security bound however is not tight.

Theorem A.7.0.2. *For any quantum adversary $\mathcal{A}$ that makes at most q_H queries to the quantum random oracle H, at most q_G queries to the quantum random oracle G, and at most q_D (classical) queries to the decapsulation oracle,*

---

[3]The security notion of OW-PCA, with access to a ciphertext $\underline{V}$alidity checking oracle.

Table 29: IND-CPA games for r5_cpa_kem: Games G_6, G_7 and G_8

Game G_6	Game G_7	Game G_8
1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$, $\boldsymbol{B_q} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times \overline{n}}$.	1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$, $\boldsymbol{B_q} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times \overline{n}}$.	1. $\boldsymbol{A} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times d/n}$, $\boldsymbol{B_q} \xleftarrow{\$} \mathcal{R}_{n,q}^{d/n \times \overline{n}}$.
2. Choose $b \xleftarrow{\$} \{0,1\}$.	2. Choose $b \xleftarrow{\$} \{0,1\}$.	2. Choose $b \xleftarrow{\$} \{0,1\}$.
3. $m \xleftarrow{\$} \{0,1\}^{\kappa}, \boldsymbol{\zeta} = ECC\_Enc_{\kappa,f}(m)$	3. $m \xleftarrow{\$} \{0,1\}^{\kappa}, \boldsymbol{\zeta} = ECC\_Enc_{\kappa,f}(m)$	3. $m \xleftarrow{\$} \{0,1\}^{\kappa}, \boldsymbol{\zeta} = ECC\_Enc_{\kappa,f}(m)$
4. $\boldsymbol{R} \xleftarrow{\$} \left(\mathcal{H}_{n,d/n}(h) \right)^{1 \times \overline{m}}$.	4. -	4. -
5. $\boldsymbol{U}' = \mathrm{R}_{q \to z, h_2}(\langle \boldsymbol{A}^T \boldsymbol{R} \rangle_{\Phi_{n+1}})$, $z = \max(p, tq/p)$	5. $\boldsymbol{U}" \xleftarrow{\$} \mathcal{R}_{n,z}^{d/n \times \overline{m}}$	5. $\boldsymbol{U}" \xleftarrow{\$} \mathcal{R}_{n,z}^{d/n \times \overline{m}}$
6. $\boldsymbol{W}" = \mathrm{R}_{q \to z, h_2}(\langle \boldsymbol{B_q}^T \boldsymbol{R} \rangle_{\xi})$	6. $\boldsymbol{W}''' \xleftarrow{\$} R_{n,z}^{\overline{n} \times \overline{m}}$	6. $\boldsymbol{W}''' \xleftarrow{\$} R_{n,z}^{\overline{n} \times \overline{m}}$
7. $\boldsymbol{v}" = \langle \mathrm{Sample}_{\mu}(\boldsymbol{W}") + \frac{pz}{qb}\boldsymbol{\zeta} \rangle_z$	7. $\boldsymbol{v}''' = \langle \mathrm{Sample}_{\mu}(\boldsymbol{W}''') + \frac{pz}{qb}\boldsymbol{\zeta} \rangle_z$	7. $\boldsymbol{v}''' = \langle \mathrm{Sample}_{\mu}(\boldsymbol{W}''') + \frac{pz}{qb}\boldsymbol{\zeta} \rangle_z$
8. $K_0 = H\left(m, \mathtt{bin}_1 \left(ct = \left(\lfloor \frac{p}{z}\boldsymbol{U}' \rfloor, \langle \lfloor \frac{tq}{pz}\boldsymbol{v}" \rfloor \rangle_t \right) \right) \right)$.	8. $K_0 = H\left(m, \mathtt{bin}_1 \left(ct = \left(\lfloor \frac{p}{z}\boldsymbol{U}" \rfloor, \langle \lfloor \frac{tq}{pz}\boldsymbol{v}''' \rfloor \rangle_t \right) \right) \right)$.	8. $K_0 \xleftarrow{\$} \{0,1\}^{\kappa}$
9. $K_1 \xleftarrow{\$} \{0,1\}^{\kappa}$.	9. $K_1 \xleftarrow{\$} \{0,1\}^{\kappa}$.	9. $K_1 \xleftarrow{\$} \{0,1\}^{\kappa}$.
10. $b' = \mathcal{A}\left(\left((\boldsymbol{A}, \langle \boldsymbol{B_q} \rangle_p), (\lfloor \frac{p}{z}\boldsymbol{U}' \rfloor, \langle \lfloor \frac{tq}{pz}\boldsymbol{v}" \rfloor \rangle_t), K_b \right) \right)$.	10. $b' = \mathcal{A}(\left((\boldsymbol{A}, \langle \boldsymbol{B_q} \rangle_p), (\lfloor \frac{p}{z}\boldsymbol{U}" \rfloor, \langle \lfloor \frac{tq}{pz}\boldsymbol{v}''' \rfloor \rangle_t), K_b \right))$.	10. $b' = \mathcal{A}\left(\left((\boldsymbol{A}, \langle \boldsymbol{B_q} \rangle_p), (\lfloor \frac{p}{z}\boldsymbol{U}" \rfloor, \langle \lfloor \frac{tq}{pz}\boldsymbol{v}''' \rfloor \rangle_t), K_b \right) \right)$.
11. Output $[b' = b]$.	11. Output $[(b' = b)]$.	10. Output $[(b' = b)]$.

there exists a quantum adversary $\mathcal{B}$ such that

$$Adv_{CCA\text{-}KEM}^{IND\text{-}CCA}(\mathcal{A}) \le 4q_H \sqrt{q_D \cdot q_H \cdot \delta + q_G \cdot \sqrt{Adv_{CPA\text{-}PKE}^{IND\text{-}CPA}(\mathcal{B})}} \qquad (61)$$

A.8   Hardness of Sparse-Ternary LWR

In this section, we prove the hardness of the Decision-LWR problem with sparse-ternary secrets assuming that the small modulus p divides the large modulus q. Figure 4 provides an overview of the reductions involved in the proof of the main result, Theorem A.8.0.1.

Theorem A.8.0.1. *Let $k, p, q \ge 1$ and $m \ge n \ge h \ge 1$ be integers such that p divides q, and $k \ge m' = \frac{q}{p} \cdot m$. Let $\epsilon \in (0, \frac{1}{2})$, and $\alpha, \delta > 0$ such that*

$$\alpha \ge q^{-1}\sqrt{(2/\pi)\ln(2n(1+\epsilon^{-1}))}, \quad \binom{n}{h}2^h \ge q^{k+1} \cdot \delta^{-2}, \text{ and } m = O(\frac{\log n}{\alpha\sqrt{10h}})$$

There exist three (transformation) reductions from $\mathtt{dLWE}_{k,m',q,D_{\alpha}}$ to $\mathtt{dLWE}_{n,m',q,D_{\alpha\sqrt{10h}}}(\mathcal{U}(\mathcal{H}_n(h)))$ such that for any algorithm for the latter problem with advantage ζ, at least one of the reductions produces an algorithm for the former with advantage at least

$$(\zeta - \delta)/(3m') - 41\epsilon/2 - \sum_{s|q, s \text{ prime}} s^{-k-1}.$$
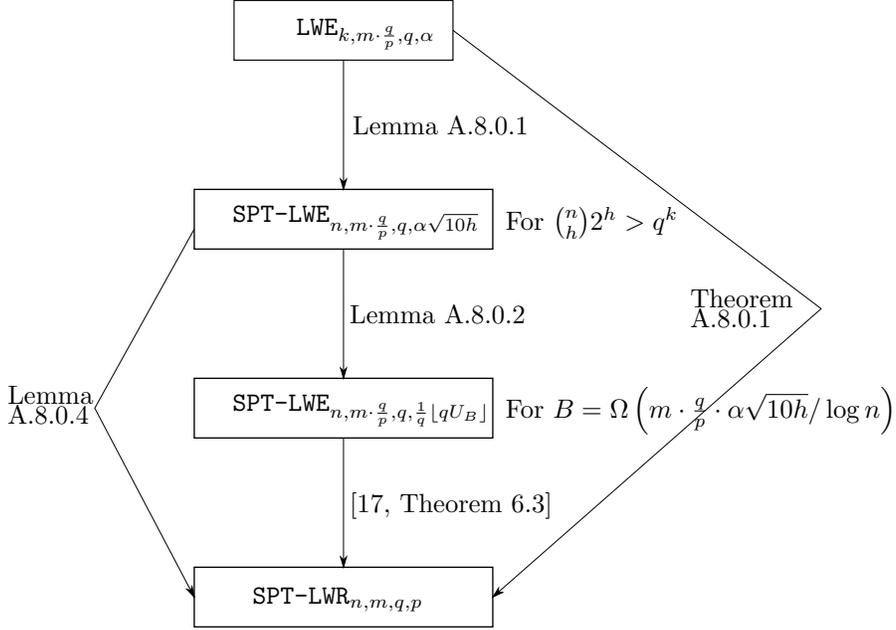
Figure 4: Summary of reductions used in Theorem A.8.0.1. "SPT" refers to a variant of the problem in question where the secret is sparse-ternary, instead of uniform in $\mathbb{Z}_q^d$.

Moreover, there is a reduction from $\mathtt{dLWE}_{n,m',q,D_{\alpha\sqrt{10h}}}\left(\mathcal{U}(\mathcal{H}_n\left(h\right))\right)$ to $\mathtt{dLWR}_{n,m,q,p}\left(\mathcal{U}(\mathcal{H}_n\left(h\right))\right)$.

Proof. Combination of Lemma A.8.0.1 and Lemma A.8.0.4 with $\alpha' = \alpha\sqrt{10h}$. $\qquad\square$

Theorem A.8.0.1 implies the hardness of the sparse-ternary LWR problem $\mathtt{LWR}_{\mathrm{spt}}$ based on the hardness of the $\mathtt{LWE}$ problem with uniformly random secrets in $\mathbb{Z}_q$ and Gaussian errors.

Step 1: Reduction from LWE with secrets in $\mathbb{Z}_q$ and Gaussian errors to Sparse-ternary LWE: In [35, Theorem 1], specializing [34, Theorem 4], it is shown that if $\binom{n}{h}2^h > q^{k+1}$ and $\omega > \alpha\sqrt{10h}$, then the $\mathtt{dLWE}_{n,m,q,D_\omega}\left(\mathcal{U}(\mathcal{H}_n\left(h\right))\right)$ problem is at least as hard as the $\mathtt{dLWE}_{k,m,q,D_\alpha}$ problem. More formally, generalizing [31, Theorem 4.1], the following holds.

Lemma A.8.0.1. *Let $k, q \geq 1$ and $m \geq n \geq h \geq 1$ be integers, and let $\epsilon \in (0, \frac{1}{2})$, and $\alpha, \delta > 0$ such that*

$$\alpha \geq q^{-1}\sqrt{(2/\pi)\ln(2n(1+\epsilon^{-1}))}, \ \text{and} \ \binom{n}{h}2^h \geq q^{k+1} \cdot \delta^{-2}$$

There exist three (transformation) reductions from $\mathtt{dLWE}_{k,m,q,D_\alpha}$ to $\mathtt{dLWE}_{n,m,q,D_{\alpha\sqrt{10h}}}(\mathcal{U}(\mathcal{H}_n(h)))$ such that for any algorithm for the latter problem with advantage ζ, at least one of the reductions produces an algorithm for the former with advantage at least

$$(\zeta - \delta)/(3m) - 41\epsilon/2 - \sum_{s|q,s\ prime} s^{-k-1}.$$

Step 2: Reduction from Sparse-ternary LWE to Sparse-ternary LWR:
Bai et al. provide in [17, Theorem 6.4] a reduction from LWE with Gaussian noise to LWR, that is based on two independent reductions It can readily be seen that one of these reductions [17, Theorem 6.3] holds for any secret distribution with support on $\mathbb{Z}_q^{n*} = \{(x_1,\ldots,x_n) \in \mathbb{Z}_q^n \mid \gcd(x_1,x_2,\ldots,x_n,q) = 1\}$, and therefore can be applied to the case when the secret is chosen from $\{-1,0,1\}^n$. The other reduction [17, Theorem 5.1] however, implicitly assumes the secret to be chosen uniformly at random from $\mathbb{Z}_q^n$. Below, we describe an extension of [17, Theorem 5.1] that describes a reduction from LWE with Gaussian noise and sparse ternary secrets reduces to LWR with sparse-ternary secrets. Below, we will describe such an extension. U_B denotes the continuous uniform distribution in $[-B,\ldots,B]$.

Lemma A.8.0.2 (Adapted from [17, Theorem 5.1])**.** *Let n, m, q be positive integers. Let $\alpha, B > 0$ be real numbers with $B = \Omega\left(m\alpha/\log n\right)$ and $Bq \in \mathbb{Z}$. Let $m > \log\left(\binom{n}{h}2^h\right)/\log\left(\alpha + B\right)^{-1} \geq 1$. Then there is a polynomial time reduction from $\mathtt{LWE}_{n,m,q,D_\alpha}(\mathcal{U}(\mathcal{H}_n(h)))$ to $\mathtt{LWE}_{n,m,q,\phi}(\mathcal{U}(\mathcal{H}_n(h)))$ with $\phi = \frac{1}{q}\lfloor qU_B \rceil$.*

Proof. The reduction proceeds similar to that of [17, Theorem 5.1], relying on five steps. Steps 1, 3, 4 below proceed exactly as in [17, Theorem 5.1]. For steps 2 and 5, we mention our adaptations in order to prove the reduction for the case of sparse-ternary secrets, and the resulting conditions. We omit details for brevity.

1. A reduction from $\mathtt{dLWE}_{n,m,q,D_\alpha}$ to $\mathtt{dLWE}_{n,m,q,\psi}$, with $\psi = D_\alpha + U_B$.

2. A reduction from $\mathtt{dLWE}_{n,m,q,\psi}$ to $\mathtt{sLWE}_{n,m,q,\psi}$. We adapt the corresponding step in [17, Theorem 5.1] to work for the uniform distribution on $\mathcal{H}_n(h)$ instead of the uniform distribution on $\mathbb{Z}_q^n$. This results in the bound on m as stated in the lemma.

3. A reduction from $\mathtt{sLWE}_{n,m,q,\psi}$ to $\mathtt{sLWE}_{n,m,q,U_B}$.

4. A reduction from $\mathtt{sLWE}_{n,m,q,U_B}$ to $\mathtt{sLWE}_{n,m,q,\phi}$, with $\phi = \frac{1}{q}\lfloor qU_B \rceil$.

5. A reduction from $\mathtt{sLWE}_{n,m,q,\phi}$ to $\mathtt{dLWE}_{n,m,q,\phi}$. Since the modulus q is not a prime, the argument from [17, Theorem 5.1] cannot be applied. Instead, we extend an argument due to Regev (see, e.g, [83]) to prove the search-to-decision reduction, which requires that Bq is an integer. We first state an easy lemma.

Lemma A.8.0.3. *Let $a > 1$, and let ϕ be the discrete probability distribution obtained by rounding the continuous uniform probability on $[-a, a]$ to the closest integer. If a is an integer, then $\sum_{k\,even}\phi(k) = \sum_{k\,odd}\phi(k) = \frac{1}{2}$.*

Proof. For $|k| \leq \lfloor a \rfloor - 1$, the interval $[k - \frac{1}{2}, k + \frac{1}{2}]$ is a subset of $-[a, a]$, so that $\sum_{k \equiv 1 - \lfloor a \rfloor \pmod 2} \phi(k) = \sum_{j=0}^{\lfloor a \rfloor - 1} \phi(2j - \lfloor a \rfloor + 1) = \frac{\lfloor a \rfloor}{2a}$. $\square$

We are now in a position to extend Regev's reduction. Let ϕ be a probability distribution on $\mathbb{Z}_q$ such that $\sum_k \phi(2k) = \sum_k \phi(2k+1) = \frac{1}{2}$. For each $\mathbf{s} \in \mathbb{Z}_q^n$, the probability distribution $A_{\mathbf{s},\phi}$ on $\mathbb{Z}_q^n \times \mathbb{Z}_q$ is obtained by choosing $\mathbf{a} \in \mathbb{Z}_q^n$ uniformly, choosing e according to ϕ, and outputting $(\mathbf{a}, (\mathbf{a}, \mathbf{s}) + e)$ where the additions are modulo q. If qB is integer, then a distinguisher for $\mathtt{dLWE}_{n,m,q,\phi}(D_s)$ can be used to construct a solver for $\mathtt{sLWE}_{n,m,q,\phi}(D_s)$ for any secret distribution D_s supported on $\{-1, 0, 1\}^n$, where ϕ is the discrete noise $\frac{1}{q}\lfloor qU_B \rfloor$. Note that if Bq is integer, the noise ϕ is distributed as $\phi(k) = \frac{1}{2B}$ for $|k| \leq B - 1$, and $\phi(B) = \phi(-B) = \frac{1}{4B}$.

We now show that if Bq is integer, then a distinguisher for deciding between uniform samples $(\mathbf{a}, u) \in U(\mathbb{Z}_q^n) \times U(Z_q)$ and samples $(\mathbf{a}, b)$ from $A_{\mathbf{s},\phi}$ for some unknown $s \in \mathcal{S} \subset \{-1, 0, 1\}^n$ can be used for solving. We show how to find s_1, the first coordinate of a secret. For each $k \in \mathbb{Z}_q$, we consider the following transformation. For each pair $(\mathbf{a}, b)$, we choose a random $r \in \mathbb{Z}_q$ and output $(\mathbf{a}', b') = (\mathbf{a} + (r, 0, \ldots, 0), b + rk)$. Clearly, this transformation takes the uniform distribution to itself. So let us now assume that $b = (\mathbf{a}, \mathbf{s}) + e$ for some $s \in \mathcal{S}$ and some error e. Then $b' = (\mathbf{a}', \mathbf{s}) + r(k - s_1) + e$. If $k = s_1$, then $(\mathbf{a}', b')$ is from $A_{\mathbf{s},\phi}$. If $|k - s_1| = 1$, then $r(k - s_1)$ is uniform over $\mathbb{Z}_q$, and so $(\mathbf{a}', \mathbf{b})$ follows the uniform distribution. Finally, we can have that $|k - s_1| = 2$. We consider $k - s_1 = 2$, the other case being similar. We then have that $b' = (\mathbf{a}, \mathbf{s}) + 2r + e \pmod q$. If q is odd, then $2r$ is uniformly distributed on $\mathbb{Z}_q$, so that $(\mathbf{a}', \mathbf{b})$ is uniformly distributed. If q is even, then $2r$ is distributed uniformly on the even elements of $\mathbb{Z}_q$. With our specific error distribution, e is even with probability $\frac{1}{2}$, so that $2r + e$ is distributed uniformly on $\mathbb{Z}_q$. So also in this case, $(\mathbf{a}', b)$ is distributed uniformly.

     $\square$

Finally, we state the reduction from $\mathtt{dLWE}_{n,m,q,D_\alpha}$ to $\mathtt{dLWR}_{n,m,q,p}$, for the sparse-ternary secret distribution.

Lemma A.8.0.4. *Let p, q be positive integers such that p divides q. Let $\alpha' > 0$. Let $m' = m \cdot (q/p)$ with $m = O(\log n/\alpha')$ for $m' \geq m \geq n \geq 1$. There is a polynomial time reduction from $\mathtt{dLWE}_{n,m',q,D_{\alpha'}}$ to $\mathtt{dLWR}_{n,m,q,p}$, both defined for the sparse-ternary secret distribution.*

Proof. Let $B = q/2p$. The reduction has two steps:

1. A reduction from $\mathtt{dLWE}_{n,m',q,D_{\alpha'}}$ to $\mathtt{dLWE}_{n,m',q,\phi}$, where $B = \Omega(m'\alpha'/\log n)$. due to Lemma A.8.0.2.

2. A reduction from $\mathtt{dLWE}_{n,m',q,\phi}$ to $\mathtt{dLWR}_{n,m,q,p}$, due to [17, Theorem 6.3].

As $m' = m \cdot (q/p) = (q/p)O(\frac{\log n}{\alpha'})$, it follows that $B = q/2p = \Omega(m'\alpha'/\log n)$, so that Lemma A.8.0.2 indeed is applicable. □

Note that the conditions imposed by Lemma]A.8.0.2 imply that $1/\alpha$ must at least grow linearly in n. This is a common bottleneck in all known LWE to LWR reductions [17, 23, 18].

References

[1] Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. Solving the shortest vector problem in 2^n time via discrete Gaussian sampling. In *STOC*, pages 733–742, 2015.

[2] Miklós Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions. In *STOC*, pages 10–19, 1998.

[3] Miklós Ajtai, Ravi Kumar, and Dandapani Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610, 2001.

[4] Martin Albrecht, Léo Ducas, Gottfried Herold, Elena Kirshanova, Eamonn Postlethwaite, and Marc Stevens. The general sieve kernel and new records in lattice reduction. *Preprint*, 2018.

[5] Martin R. Albrecht. On dual lattice attacks against small-secret LWE and parameter choices in HElib and SEAL. Cryptology ePrint Archive, Report 2017/047, 2017. `http://eprint.iacr.org/2017/047`.

[6] Martin R. Albrecht. On dual lattice attacks against small-secret LWE and parameter choices in HElib and SEAL. Cryptology ePrint Archive, Report 2017/047, 2017.

[7] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - a new hope. Cryptology ePrint Archive, Report 2015/1092, 2015. `http://eprint.iacr.org/2015/1092`.

[8] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - a new hope. Cryptology ePrint Archive, Report 2015/1092, 2015.

[9] Joel Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. Learning with rounding, revisited: New reduction, properties and applications. Cryptology ePrint Archive, Report 2013/098, 2013. `http://eprint.iacr.org/2013/098`.

[10] Alexandr Andoni, Thijs Laarhoven, Ilya Razenshteyn, and Erik Waingarten. Optimal hashing-based time-space trade-offs for approximate near neighbors. In *SODA*, pages 47–66, 2017.

[11] Yoshinori Aono and Phong Q. Nguyên. Random sampling revisited: lattice enumeration with discrete pruning. In *EUROCRYPT*, 2017.

[12] Yoshinori Aono, Phong Q. Nguyen, and Yixin Shen. Quantum lattice enumeration and tweaking discrete pruning. In *ASIACRYPT*, 2018.

[13] Hayo Baan, Sauvik Bhattacharya, Oscar Garcia-Morchon, Ronald Rietman, Ludo Tolhuizen, Jose-Luis Torre-Arce, and Zhenfei Zhang. Round2: KEM and PKE based on GLWR, November 2017.

[14] Hayo Baan, Sauvik Bhattacharya, Oscar Garcia-Morchon, Ronald Rietman, Ludo Tolhuizen, Jose-Luis Torre-Arce, and Zhenfei Zhang. Round2: KEM and PKE based on GLWR. Cryptology ePrint Archive, Report 2017/1183, 2017.

[15] L. Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.

[16] Shi Bai and Steven D. Galbraith. Lattice Decoding Attacks on Binary LWE. Cryptology ePrint Archive, Report 2013/839, 2013. `http://eprint.iacr.org/2013/839`.

[17] Shi Bai, Adeline Langlois, Tancrède Lepoint, Amin Sakzad, Damien Stehle, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: using the Rényi divergence rather than the statistical distance. Cryptology ePrint Archive, Report 2015/483, 2015. `http://eprint.iacr.org/2015/483`.

[18] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. Cryptology ePrint Archive, Report 2011/401, 2011. `http://eprint.iacr.org/2011/401`.

[19] Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. In *SODA*, pages 10–24, 2016.

[20] Anja Becker and Thijs Laarhoven. Efficient (ideal) lattice sieving using cross-polytope LSH. In *AFRICACRYPT*, pages 3–23, 2016.

[21] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, CCS '93, pages 62–73, New York, NY, USA, 1993. ACM.

[22] Sauvik Bhattacharya, Oscar Garcia-Morchon, Ronald Rietman, and Ludo Tolhuizen. spKEX: An optimized lattice-based key exchange. Cryptology ePrint Archive, Report 2017/709, 2017.

[23] Andrej Bogdanov, Siyao Guo, Daniel Masny, Silas Richelson, and Alon Rosen. On the Hardness of Learning with Rounding over Small Modulus. Cryptology ePrint Archive, Report 2015/769, 2015. `http://eprint.iacr.org/2015/769`.

[24] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. Cryptology ePrint Archive, Report 2010/428, 2010. `http://eprint.iacr.org/2010/428`.

[25] Joppe Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! Practical, Quantum-Secure Key Exchange from LWE. Cryptology ePrint Archive, Report 2016/659, 2016. `http://eprint.iacr.org/2016/659`.

[26] Joppe Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, and Damien Stehlé. CRYSTALS – Kyber: a CCA-secure module-lattice-based KEM. Cryptology ePrint Archive, Report 2017/634, 2017. `http://eprint.iacr.org/2017/634`.

[27] Joppe Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, and Damien Stehlé. CRYSTALS – Kyber: a CCA-secure module-lattice-based KEM. Cryptology ePrint Archive, Report 2017/634, 2017.

[28] Joppe W. Bos, Craig Costello, Michael Naehrig, and Douglas Stebila. Post-Quantum Key Exchange for the TLS Protocol from the Ring Learning with Errors Problem. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 553–570, 2015.

[29] Joppe W. Bos, Michael Naehrig, and Joop van de Pol. Sieving for shortest vectors in ideal lattices: a practical perspective. Cryptology ePrint Archive, Report 2014/880, 2014.

[30] Joppe W. Bos, Michael Naehrig, and Joop van de Pol. Sieving for shortest vectors in ideal lattices: a practical perspective. *International Journal of Applied Cryptography*, pages 1–23, 2016.

[31] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical Hardness of Learning with Errors. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 575–584, New York, NY, USA, 2013. ACM.

[32] Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better lattice security estimates. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASiACRYPT 2011*, volume 7073 of *LNCS*, pages 1–20. Springer, 2011.

[33] Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better Lattice Security Estimates. In *Proceedings of the 17th International Conference on The Theory and Application of Cryptology and Information Security*, ASIACRYPT'11, pages 1–20, Berlin, Heidelberg, 2011. Springer-Verlag.

[34] Jung Hee Cheon, Kyoo Hyung Han, Jinsu Kim, Changmin Lee, and Yongha Son. A Practical Post-Quantum Public-Key Cryptosystem Based on spLWE. Cryptology ePrint Archive, Report 2016/1055, 2016. `http://eprint.iacr.org/2016/1055`.

[35] Jung Hee Cheon, Duhyeong Kim, Joohee Lee, and Yongsoo Song. Lizard: Cut off the Tail! Practical Post-Quantum Public-Key Encryption from LWE and LWR. Cryptology ePrint Archive, Report 2016/1126, 2016. http://eprint.iacr.org/2016/1126.

[36] John H. Conway and Neil J.A. Sloane. *Sphere packings, lattices and groups.* Springer, 1999.

[37] Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård Revisited: How to Construct a Hash Function. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, pages 430–448. Springer, 2005.

[38] Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. Cryptology ePrint Archive, Report 2001/108, 2001.

[39] Ronald Cramer and Victor Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. Cryptology ePrint Archive, Report 2001/108, 2001. http://eprint.iacr.org/2001/108.

[40] Jan-Pieter D'Anvers, Angshuman Karmakar, Sujoy Sinha Roy, and Frederik Vercauteren. Saber: Module-LWR based key exchange, CPA-secure encryption and CCA-secure KEM. Cryptology ePrint Archive, Report 2018/230, 2018.

[41] TU Darmstadt. SVP challenge, 2018.

[42] Léo Ducas. Shortest vector from lattice sieving: a few dimensions for free. In *EUROCRYPT*, 2018.

[43] Morris J Dworkin. Sha-3 standard: Permutation-based hash and extendable-output functions. Technical report, National Institute for Standards and Technology, 2015.

[44] FIPS. Specification for the Advanced Encryption Standard (AES). Federal Information Processing Standards Publication 197, November 2001.

[45] FIPS. SHA-3 standard: Permutation-based hash and extendable-output functions. Federal Information Processing Standards Publication 202, August 2015.

[46] Tim Fritzmann, Thomas Pöppelmann, and Johanna Sepulveda. Analysis of error-correcting codes for lattice-based key exchange. Cryptology ePrint Archive, Report 2018/150, 2018.

[47] Nicolas Gama, Phong Q. Nguyên, and Oded Regev. Lattice enumeration using extreme pruning. In *EUROCRYPT*, pages 257–278, 2010.

[48] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 212–219. ACM, 1996.

[49] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 212–219, New York, NY, USA, 1996. ACM.

[50] Mike Hamburg. Graphs of "estimate all the LWE, NTRU schemes!" indexed to the "pqc lounge" data., 2017.

[51] Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In *CRYPTO*, pages 447–464, 2011.

[52] Gottfried Herold, Elena Kirshanova, and Thijs Laarhoven. Speed-ups and time-memory trade-offs for tuple lattice sieving. In *PKC*, pages 407–436, 2018.

[53] Dennis Hoffheinz, Kathrin Hövelmanns, and Eike Kiltz. A Modular Analysis of the Fujisaki-Okamoto Transformation. Cryptology ePrint Archive, Report 2017/604, 2017. `http://eprint.iacr.org/2017/604`.

[54] Jeffrey Hoffstein, Jill Pipher, John M. Schanck, Joseph H. Silverman, William Whyte, and Zhenfei Zhang. Choosing Parameters for NTRUEncrypt, 2017.

[55] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In Joe P. Buhler, editor, *Algorithmic Number Theory, Third International Symposium, ANTS-III, Portland, Oregon, USA, June 21-25, 1998, Proceedings*, pages 267–288. Springer, 1998.

[56] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. Cryptology ePrint Archive, Report 2017/604, 2017.

[57] Nick Howgrave-Graham. A Hybrid Lattice-Reduction and Meet-in-the-Middle Attack Against NTRU. In Alfred Menezes, editor, *Advances in Cryptology - CRYPTO 2007: 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007. Proceedings*, pages 150–169. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.

[58] Andreas Hülsing, Joost Rijneveld, John M. Schanck, and Peter Schwabe. High-speed key encapsulation from NTRU. In *Cryptographic Hardware and Embedded Systems - CHES 2017 - 19th International Conference, Taipei, Taiwan, September 25-28, 2017, Proceedings*, pages 232–252, 2017.

[59] Andreas Hülsing, Joost Rijneveld, John M. Schanck, and Peter Schwabe. High-speed key encapsulation from NTRU. In *Cryptographic Hardware and Embedded Systems - CHES 2017 - 19th International Conference, Taipei, Taiwan, September 25-28, 2017, Proceedings*, pages 232–252, 2017.

[60] Tsukasa Ishiguro, Shinsaku Kiyomoto, Yutaka Miyake, and Tsuyoshi Takagi. Parallel Gauss Sieve algorithm: Solving the SVP challenge over a 128-dimensional ideal lattice. In *PKC*, pages 411–428, 2014.

[61] Ravi Kannan. Improved algorithms for integer programming and related lattice problems. In *STOC*, pages 193–206, 1983.

[62] Subhash Khot. Hardness of approximating the shortest vector problem in lattices. In *FOCS*, pages 126–135, 2004.

[63] Aleksandr Nikolayevich Korkin and Yegor Ivanovich Zolotarev. Sur les formes quadratiques. *Mathematische Annalen*, 6(3):366–389, 1873.

[64] Thijs Laarhoven. *Search problems in cryptography.* PhD thesis, Eindhoven University of Technology, 2015.

[65] Thijs Laarhoven. *Search problems in cryptography, From fingerprinting to lattice sieving.* PhD thesis, Eindhoven University of Technology, 2016.

[66] Thijs Laarhoven and Artur Mariano. Progressive lattice sieving. In *PQCrypto*, 2018.

[67] Thijs Laarhoven, Michele Mosca, and Joop van de Pol. Finding shortest lattice vectors faster using quantum search. Cryptology ePrint Archive, Report 2014/907, 2014.

[68] Thijs Laarhoven, Michele Mosca, and Joop van de Pol. Finding shortest lattice vectors faster using quantum search. *Designs, Codes and Cryptography*, 77(2):375–400, 2015.

[69] Arjen Klaas Lenstra, Hendrik Willem Lenstra, and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.

[70] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On Ideal Lattices and Learning with Errors Over Rings. Cryptology ePrint Archive, Report 2012/230, 2012. http://eprint.iacr.org/2012/230.

[71] Ueli Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. Cryptology ePrint Archive, Report 2003/161, 2003.

[72] Daniele Micciancio and Panagiotis Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In *STOC*, pages 351–358, 2010.

[73] Daniele Micciancio and Michael Walter. Fast lattice point enumeration with minimal overhead. In *SODA*, pages 276–294, 2015.

[74] Daniele Micciancio and Michael Walter. Practical, predictable lattice basis reduction. In *EUROCRYPT*, pages 820–849, 2016.

[75] Michael Naehrig, Erdem Alkim, Joppe Bos, Leo Ducas, Karen East-erbrook, Brian LaMacchia, Patrick Longa, Ilya Mironov, Valeria Niko-laenko, Christopher Peikert, Ananth Raghunathan, and Douglas Ste-bila. FrodoKEM. Technical report, National Institute of Standards and Technology, 2017. Available at https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions.

[76] National Institute of Standards and Technology. Submission Re-quirements and Evaluation Criteria for the Post-Quantum Cryp-tography Standardization Process. POST-QUANTUM CRYPTO STANDARDIZATION. Call For Proposals Announcement, 2016. http://csrc.nist.gov/groups/ST/post-quantum-crypto/documents/call-for-proposals-final-dec-2016.pdf.

[77] Chris Peikert. Lattice cryptography for the internet. Cryptology ePrint Archive, Report 2014/070, 2014. http://eprint.iacr.org/2014/070.

[78] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudoran-domness of ring-LWE for any ring and modulus. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 461–473, 2017.

[79] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudoran-domness of Ring-LWE for Any Ring and Modulus, 2017.

[80] Michael E. Pohst. On the computation of lattice vectors of minimal length, successive minima and reduced bases with applications. *ACM SIGSAM Bulletin*, 15(1):37–44, 1981.

[81] Xavier Pujol and Damien Stehlé. Solving the shortest lattice vector problem in time $2^{2.465n}$. *Cryptology ePrint Archive, Report 2009/605*, pages 1–7, 2009.

[82] Charles Rackoff and Daniel R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In Joan Feigenbaum, editor, *Advances in Cryptology — CRYPTO '91: Proceedings*, pages 433–444. Springer, 1992.

[83] Oded Regev. On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In *Proceedings of the Thirty-seventh Annual ACM Sym-posium on Theory of Computing*, STOC '05, pages 84–93, New York, NY, USA, 2005. ACM.

[84] Oded Regev. The Learning with Errors Problem (Invited Survey). In *Pro-ceedings of the 25th Annual IEEE Conference on Computational Complex-ity, CCC 2010, Cambridge, Massachusetts, June 9-12, 2010*, pages 191–204, 2010.

[85] Markku-Juhani O. Saarinen. HILA5: Key Encapsulation Mechanism (KEM) and Public Key Encryption Algorithm, November 2017.

[86] Markku-Juhani O. Saarinen. Ring-LWE ciphertext compression and error correction: Tools for lightweight post-quantum cryptography. In *Proceedings of the 3rd ACM International Workshop on IoT Privacy, Trust, and Security*, IoTPTS '17, pages 15–22. ACM, April 2017.

[87] Markku-Juhani O. Saarinen. HILA5: On reliability, reconciliation, and error correction for Ring-LWE encryption. In Carlisle Adams and Jan Camenisch, editors, *SAC 2017*, volume 10719 of *Lecture Notes in Computer Science*, pages 192–212. Springer, 2018.

[88] Michael Schneider. Sieving for shortest vectors in ideal lattices. Cryptology ePrint Archive, Report 2011/458, 2011.

[89] C. P. Schnorr and M. Euchner. Lattice Basis Reduction: Improved Practical Algorithms and Solving Subset Sum Problems. *Math. Program.*, 66(2):181–199, September 1994.

[90] C. P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Math. Program.*, 66(2):181–199, September 1994.

[91] Claus-Peter Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science*, 53(2–3):201–224, 1987.

[92] Claus-Peter Schnorr and Horst Helmut Hörner. Attacking the Chor-Rivest cryptosystem by improved lattice reduction. In *EUROCRYPT*, pages 1–12, 1995.

[93] Naftali Sommer, Meir Feder, and Ofir Shalvi. Finding the closest lattice point by iterative slicing. *SIAM Journal of Discrete Mathematics*, 23(2):715–731, 2009.

[94] Thomas Wunderer. Revisiting the hybrid attack: Improved analysis and refined security estimates. Cryptology ePrint Archive, Report 2016/733, 2016. http://eprint.iacr.org/2016/733.

[95] Shang-Yi Yang, Po-Chun Kuo, Bo-Yin Yang, and Chen-Mou Cheng. Gauss sieve algorithm on GPUs. In *CT-RSA*, pages 39–57, 2017.