Provable decryption failure security for practical lattice-based PKE

Christian Majenz[0000-0002-1877-8385] and Fabrizio Sisinni[0009-0007-9641-4329]

Technical University of Denmark, Denmark

Abstract. Recently, Hövelmanns, Hülsing, and Majenz introduced a security notion called Find Failing Plaintext – Non Generic (FFP-NG), which captures the ability of an adversary to find decryption failures by making non-trivial use of the public key. A first analysis of this property for lattice-based schemes was presented by Majenz and Sisinni, who showed that the Learning With Errors (LWE) problem reduces to breaking the FFP-NG security of the PVW scheme with discrete Gaussian noise. In this work, we generalize their result by analysing the FFP-NG security of widely used schemes based on Ring-LWE and Module-LWE. To keep our analysis as general as possible, we consider a family of subgaussian distributions that includes, among others, discrete Gaussians and centered binomials.

Keywords: Post-Quantum Cryptography · Lattice-based Cryptography · Provable Security

Table of Contents

1 Introduction

Since its introduction [1, 33], lattice-based cryptography has emerged as a versatile foundation for a wide range of cryptographic constructions [31]. The introduction of the Learning With Errors (LWE) problem [33] marked a turning point, and subsequent algebraic variants—most notably Ring-LWE (RLWE) [26, 27] and Module-LWE (MLWE) [25]—were proposed to improve efficiency while retaining strong security guarantees. From a security standpoint, LWE and its algebraic variants are particularly attractive hardness assumptions for post-quantum cryptography: they admit worst-case to average-case reductions from lattice problems, and no known techniques exploit their algebraic structure to yield more effective quantum attacks than those applicable to standard LWE. At the same time, these variants enable public-key encryption schemes that are efficient enough for practical deployment. In particular, MLWE underlies ML-KEM [29], which is one of the two Key Encapsulation Mechanisms (KEMs) standardized by NIST and the most prominently used one in the post-quantum transition.

When balancing efficiency and security requirements for the parameters of (algebraic) LWE-based PKE schemes, it turns out to be advantageous to allow a small probability of decryption failure. This yields a PKE scheme with only **approximate** correctness. While designers keep the decryption failure probability low enough to ensure correctness in practice, decryption failures can also pose a security problem [5, 11, 14]. It is therefore important to characterize the security implications of decryption failures in LWE-based schemes.

For security proofs, the imperfect correctness of many lattice-based PKE schemes is analysed when upgrading them to IND-CCA-secure KEMs using the Fujisaki-Okamoto (FO) transformation and its variants [15, 16, 18]. For post-quantum security, these transformations are analysed in the quantum-accessible random oracle model (QROM) [6]. The standard approach, first put forward by [19] and then improved in a long line of work [4, 20–23, 34], is to replace the (imperfectly correct) decapsulation oracle with a (perfectly correct) simulation. The distinguishability of these two oracles is then bounded by reducing it to an unstructured search task. While this approach yields security proof relying on the worst-case decryption failure probability only (a statistical property), the resulting security bounds are unlikely to be tight. In [20], a more fine-grained

The authors thanks one of the anonymous reviewers of SAC 2026 for all the good feedback. The authors acknowledge support from the Danish Ministry of Defense Acquisition and Logistics Organization (FMI). CM acknowledges support by the Independent Research Fund Denmark via a DFF Sapere Aude grant (IM-3PQC, grant ID 10.46540/2064-00034B). FS and CM are part of the Quantum-Safe Internet (QSI) ITN which received funding from the European Union's Horizon-Europe program as Marie Sklodowska-Curie Action (PROJECT 101072637 - HORIZON - MSCA-2021-DN-01).

security reduction has been developed which separates statistical and computational decryption failure finding attacks. This reduction yields tighter bounds, but requires for each PKE scheme the characterization of the difficulty to find non-generic decryption failures. In the novel security game introduced for this purpose, FFP-NG ("find failing plaintexts – non-generic"), the adversary has to find decryption failures that occur more likely for a given public key than for an independently generated one. As a result, the FFP-NG security of PKE schemes needs to be characterized to exploit the tightest reduction presented in [20].

From a cryptanalysis perspective, a line of work in this direction [5, 11, 14] has culminated in a work showing that even a single decryption failure can lead to key-recovery attacks [12]. This illustrates that the above-described security losses reflect actual attacks and cannot be ignored.

There is only one previous work that proves FFP-NG security for a PKE. In [28], authors present a security reduction from the (plain) LWE problem with discrete Gaussian noise to breaking FFP-NG security of Regev's PKE scheme [33] and its variant [32] (with the same noise distribution).

1.1 Our Contributions

In this work, we generalize the security reduction developed in [28] in two directions:

1. **Extension to algebraic lattices.** We first extend their reduction to the ring setting, using as the underlying scheme the one introduced in [26], and to the module setting, analysing as the underlying PKE scheme a simplified version of the one used to build ML-KEM [9].
2. **Generalization of noise distributions.** We also broaden the class of admissible probability distributions for the (M/R)LWE noise. Specifically, we introduce a family of subgaussian distributions that encompasses, among others, continuous Gaussians, discrete Gaussians, and centered binomials. We then go on to prove that the reduction from [28] as well as its generalizations to RLWE and MLWE work for this class of noise distributions.

These two generalizations are particularly meaningful. The schemes we analyse are the natural extensions of the Dual Regev scheme [17] to rings and modules. Several protocols [2, 8, 29, 35] follow the blueprint of the Dual Regev scheme, that is, they rely on the LWE hardness assumption both for key generation and for encryption. Moreover, the class of probability distributions we allow is widely employed in both theoretical [25–27] and practical contexts, including for example ML-KEM and Newhope [2, 9, 29].

1.2 Technical overview

In the FFP-NG game for a PKE scheme Π, a challenger honestly generates two key pairs (sk_0, pk_0) and (sk_1, pk_1), samples uniformly at random a bit b, and gives the adversary always the public key pk_0 of the first key pair. The

adversary can now make a single query to a failure-checking oracle which takes as input a message-randomness pair, uses the key pair (sk_b, pk_b) to test if the input triggers a decryption failure. Using this information, the adversary wins the game if it correctly guesses the bit b used by the oracle. In other words, the adversary needs to decide whether the failure-checking oracle used the key pair corresponding to the input public key or an independent one.

For LWE-based schemes, decryption failures are independent of the message and occur with very low probability. In this setting, the only possible strategy for an FFP-NG adversary is to find a randomness value that has a higher probability of causing a decryption failure for the given public key than for an independently generated one.

To understand the challenges in the ring/module setting and when considering more general families of error distributions, we briefly describe the reduction strategy introduced in [28] in the context of plain LWE with discrete Gaussian errors. The reduction can exploit an FFP-NG adversary as follows. Given a sample $(\mathbf{A}, \mathbf{b})$ that is either an LWE sample or a uniformly random one, the reduction samples a small error $\mathbf{e}'$ and adds it to the second component of the sample in such a way that, if the sample was uniform, it remains uniform, whereas if it was an LWE sample, it remains an LWE sample but with a slightly modified error distribution. The reduction then calls the FFP-NG adversary on input the modified sample $(\mathbf{A}, \mathbf{b}')$. The intuition is that when the adversary receives a uniform pair, the randomness $\mathbf{r}$ used to query the failure-checking oracle is independent of the added error $\mathbf{e}'$, so these values are likely to be nearly orthogonal. Conversely, when the adversary receives an LWE sample and succeeds in causing a decryption failure, the randomness $\mathbf{r}$ is somewhat aligned with the error, and thus also with the added error $\mathbf{e}'$. Using this distinction, the adversary can construct a test involving the known added error and the randomness to distinguish between an LWE sample and a uniform one. To analyse the performance of this test, it is necessary to consider the distribution of the additional LWE error added by the reduction conditioned on a fixed value for the total LWE error. Fortunately, in the case of discrete Gaussian errors, this distribution is very close to a discrete Gaussian with a different mean and variance, allowing the application of standard tail bounds. In the above-described reduction strategy, it is crucial that decryption failure occurs if and only if a certain condition on the inner product (or matrix product) of LWE error and adversarial encryption randomness is fulfilled.

The picture is a bit more involved when dealing with schemes based on RLWE or MLWE. Indeed, each value is an element of a cyclotomic field or a vector of elements of a cyclotomic field. This makes the description of the decryption failure condition more intricate. There are two possible ways of describing these elements: the **canonical embedding** and the **coefficient embedding**. The former has the advantage of mapping the product of two elements to the entry-wise product of their embeddings, while the latter represents each element by the coefficient vector of its defining polynomial. Under the canonical embedding, the decryption failure condition translates into a bound on the ℓ_∞-norm of the

product of the LWE error and the adversarial randomness vector. While this embedding is quite useful when analysing the hardness assumptions and their reductions [25–27], it is not ideal in the context of proving FFP-NG security. The schemes we study are defined with the coefficient embedding in mind and operate over cyclotomic fields modulo a prime [2, 8, 9]. These sets are easily described in this representation, as they consist simply of polynomials with coefficients modulo the prime. The same is not true for the canonical embedding, and switching between the two representations introduces a loss into the reduction. For the specific family of cyclotomic fields used in practical schemes, we can exploit their algebraic structure to describe the product of elements in the coefficient embedding using a compact formula that involves again only inner products. We therefore adopt the coefficient embedding throughout.

Several theoretical results about RLWE and MLWE can be proven using a broad class of error distributions, namely **subgaussian** distributions. This class includes, among others, gaussians distributions, discrete gaussians, and all centered bounded distributions. By definition, a subgaussian random variable has mean zero and tails that can be bounded from above by gaussian tails. However, this definition provides no control over the behaviour of the distribution around its mean, which makes impossible to control the mean and tails of a conditional distribution. To circumvent this problem and generalize the reduction introduced in [28] we observe that for several relevant error distributions – such as the centered binomial distribution used by ML-KEM [29] – the conditional distribution of interest is itself a concrete subgaussian distribution with an easy-to-describe mean. This suffices to prove the effectiveness of the reduction from RLWE and MLWE to breaking FFP-NG security of the PKE schemes studied in this work. We therefore formalize this property as an explicit assumption on the error distributions.

1.3 Structure of the paper

The remainder of the paper is organized as follows:

- Sect. 2: We introduce the notation used to describe both Ring-LWE and Module-LWE. We then review relevant background from algebraic number theory, including the definitions of cyclotomic number fields and their rings of integers, and describe the coefficient embedding. We also establish several results concerning subgaussian probability distributions. Finally, we formally define the Ring-LWE and Module-LWE problems and state the corresponding worst-case to average-case reductions. We formally introduce the FFP-NG security game and discuss some basic properties of this security notion.
- Sect. 3: We present the LPR PKE scheme, discuss its correctness and IND-CPA security, and then focus on the security reduction from RLWE to FFP-NG using the LPR scheme as underlying PKE.
- Sect. 4: We describe a simplified version of the PKE scheme underlying ML-KEM and use it to extend our reduction from MLWE to FFP-NG.

2 Preliminaries

2.1 Notations

We use $\log(n)$ to denote the logarithm of n in base 2. We use lower case letters for ring elements and polynomials, upper case letters for vectors of ring elements, bold lower case letters for vectors over $\mathbb{R}^n$, and bold upper case letters for matrices. In case of matrices, it will be clear from the context if they have entries in $\mathbb{R}$ or in a generic ring $\mathcal{R}$. Given an n-dimensional vector $\mathbf{x}$, we denote with $\mathbf{x}[k]$ its k-th coordinate. We denote by $\mathbf{x}^\top$ the transpose of $\mathbf{x}$. For a vector $\mathbf{x}$ in $\mathbb{R}$ we denote the ℓ_2 norm as $\|\mathbf{x}\|$ and the ℓ_∞ norm as $\|\mathbf{x}\|_\infty$.
Given a polynomial $p(x) = \sum_{k=0}^{n-1} p_k x^k$, we denote with $\mathbf{p} = [p_0, \dots, p_{n-1}]^\top$ its coefficient vector.
For any probability distribution $\mathcal{X}$, we write $X \leftarrow \mathcal{X}$ to denote that the random variable X is sampled according to $\mathcal{X}$. Given a set S we write $X \leftarrow_\$ S$ to denote that X is sampled uniformly at random from the set S. We say that a random variable is centered if it has mean zero, and that it is t-bounded if its ℓ_∞ norm is bounded by t.
For a real number $x \in \mathbb{R}$, let $\lfloor x \rceil$ denote the integer closest to x, with ties rounded up. Furthermore, given integers p, q, we define the function $\lfloor \cdot \rceil_{p\to q} : \mathbb{Z}_p \to \mathbb{Z}_q$ by $\lfloor x \rceil_{p\to q} := \lfloor x \cdot q/p \rceil$. This function can be extended to polynomials and vectors applying the function coordinate-wise. Given a function $f(n)$, we say that f is negligible in n if $\lim_{n\to\infty} n^c \cdot f(n) = 0$, for all $c > 0$. In this case we write $f \in negl(n)$. We say that an event A occurs with overwhelming probability if $\Pr[A] = 1 - negl(n)$.

2.2 Cyclotomic Number Fields

Here we introduce the necessary background from algebraic number theory. Further details and proofs can be found in introductory books on the subject, e.g. [24].

Given an integer $m > 0$, an element $\zeta \in \mathbb{C}$ is an mth root of unity if $\zeta^m = 1$, i.e., if it is a root of the polynomial $X^m - 1$. We say that ζ is a primitive mth root of unity if $\zeta^j \neq 1$ for all $j \in \{1, \dots, m-1\}$. Let Φ_m denote the mth cyclotomic polynomial, that is, the monic irreducible polynomial over $\mathbb{Q}$ such that $\Phi_m(\zeta) = 0$ for every primitive mth root of unity $\zeta \in \mathbb{C}$. It is well known that Φ_m has degree $n = \varphi(m)$, where φ denotes Euler's totient function. When m is a power of 2, we have $n = \varphi(m) = m/2$ and $\Phi_m(X) = X^n + 1$. throughout the rest of the paper, unless stated otherwise, we assume that m is a power of 2 and ζ is a primitive mth root of unity.

The mth cyclotomic field is the field extension $\mathbb{K}_m = \mathbb{Q}(\zeta)$. An element $a \in \mathbb{K}_m$ is called an algebraic integer if and only if its minimal polynomial over $\mathbb{Q}$ has integer coefficients. It can be shown that the set of algebraic integers of $\mathbb{K}_m$ is a subring, usually denoted by $\mathcal{O}_{\mathbb{K}_m}$, and it holds $\mathcal{O}_{\mathbb{K}_m} = \mathbb{Z}[\zeta]$. Since the underlying cyclotomic field will always be clear from the context, we drop the subscript and simply write $\mathcal{R}$ for the ring of integers.

The field $\mathbb{K}_m$ is isomorphic to the quotient field $\mathbb{Q}[X]/\langle \Phi_m \rangle$. This isomorphism can be described explicitly by mapping a polynomial $f(X) \in \mathbb{Q}[X]/\langle \Phi_m \rangle$ to its evaluation $f(\zeta) \in \mathbb{K}_m$. Consequently, elements of $\mathbb{K}_m$ can be represented as polynomials over $\mathbb{Q}$ of degree at most $n-1$. Restricting this map to $\mathcal{R}$ yields an isomorphism with $\mathbb{Z}[X]/\langle \Phi_m \rangle$. Thus, elements of $\mathcal{R}$ can be represented by polynomials of degree at most $n-1$ with integer coefficients.

This representation induces an embedding from $\mathbb{K}_m$ into $\mathbb{Q}^n \subset \mathbb{R}^n$. Given an element $a \in \mathbb{K}_m$, let $a(x) = \sum_{i=0}^{n-1} a_i x^i$ denote its polynomial representation. We can define its coefficient vector as $\mathbf{a}^\top := [a_0, \ldots, a_{n-1}]\, \mathbb{Q}^n \in \mathbb{R}^n$. The function that associates a with $\mathbf{a}$ is called coefficient embedding. With a slight abuse of notation, we will often identify elements of $\mathbb{K}_m$ with their coefficient vectors. Restricting this embedding to $\mathcal{R}$ yields vectors in $\mathbb{Z}^n$. We use the embedding to define the ℓ_2 and ℓ_∞ norm over $\mathbb{K}_m$. For $a \in \mathbb{K}_m$, we define $\|a\| := \|\mathbf{a}\|$ and $\|a\|_\infty := \|\mathbf{a}\|_\infty$.

Using a notation similar to that introduced in [13], we describe the coefficient vector of the product of two elements with a compact formula. Let $a \in \mathcal{R}$, and let $\mathbf{a} \in \mathbb{Z}^n$ denote its coefficient vector. For every $k \in \{0, \ldots, n-1\}$, we define the kth rotation vector of a as the coefficient vector of the polynomial $X^k a\left(X^{-1}\right)$ mod $X^n + 1$. We denote this vector with $\mathbf{a}_{(k)}$. For every k, it holds

$$\|\mathbf{a}_{(k)}\| = \|\mathbf{a}\| \text{ and } \|\mathbf{a}_{(k)}\|_\infty = \|\mathbf{a}\|_\infty.$$

Given a, $b \in \mathcal{R}$, and let $c = ab \in \mathcal{R}$. For every $k \in \{0, \ldots, n-1\}$, we have that

$$\mathbf{c}[k] = \langle \mathbf{a}, \mathbf{b}_{(k)} \rangle. \tag{2.1}$$

We extend the definition of kth rotation to vectors of polynomials as follows. Let $A^\top = (a_1, \ldots, a_r) \in \mathcal{R}^r$, we define its kth rotation as $\mathbf{A}^\top_{(k)} := \left(\mathbf{a_1}_{(k)} || \ldots || \mathbf{a_r}_{(k)}\right) \in \mathbb{R}^{nr}$, that is, the concatenation of the kth rotation vectors of its coordinates. This formula will be useful when describing the decryption failure condition of the PKE schemes we analyse.

There is another way to embed $\mathbb{K}_m$ and its ring of integers into a real vector space, known as the canonical embedding. This embedding is a widely used tool in algebraic number theory and provides $\mathbb{K}_m$ with a natural geometric structure. It was used to introduce both the Ring-LWE problem [26, 27] and the Module-LWE problem [25], as well as to establish their connections to hard lattice problems. In addition to being a well-studied theoretical tool, one attractive feature of this embedding is that it allows one to describe both additions and multiplications in $\mathcal{R}$ in a straightforward way: in particular, multiplication in $\mathbb{K}_m$ corresponds to the coordinate-wise multiplication of embedded vectors.

However, the schemes we study are defined over cyclotomic fields modulo a prime [2, 8, 9]. These sets are efficiently described using the coefficient embedding, whereas their representation in the canonical embedding is less direct and requires switching between embeddings. This procedure introduces losses in the reduction. Since our goal is to analyse PKE schemes that are as close as possible to real implementations, we therefore adopt the coefficient embedding throughout our analysis.

2.3 Subgaussian random variables

Here we review some simple properties of subgaussian probability distributions and we introduce the subclass of subgaussian distribution we use as error distributions in this paper.

For any $\sigma > 0$, we say that a probability distribution $\mathcal{X}$ over $\mathbb{R}$ is subgaussian with parameter σ, if $X \leftarrow \mathcal{X}$ is centered and for every $t \in \mathbb{R}$ it holds

$$\mathbb{E}\left[\exp\left(2\pi t X\right)\right] \leq \exp\left(\pi\sigma^2 t^2\right). \tag{2.2}$$

With a small abuse of notation sometime we call a random variable subgaussian, meaning that it has been sampled from a subgaussian probability distribution. To relax the notation, we write $\mathcal{X}_\sigma$ is subgaussian instead of writing it is a subgaussian probability distribution with parameter σ. In case we don't specify the parameter, it means that can be deduced from the context or other properties of the distribution. In next sections we will use different subgaussian distributions, when multiple random variables are sampled from $\mathcal{X}_\sigma$ we mean that they are sampled from the same probability distribution that is subgaussian with parameter σ. Instead, when we sample from subgaussians with different parameters, we allow these random variables to have completely different distributions, not only the same distribution but different parameters.

Proposition 2.1 (Extended Markov's inequality) *Let X is a random variable over $\mathbb{R}$ and $f : \mathbb{R} \to \mathbb{R}$. If f is a non-decreasing, non-negative function, and $f(c) > 0$, then*

$$\Pr\left[X \geq c\right] \leq \frac{\mathbb{E}\left[f(X)\right]}{f(c)}. \tag{2.3}$$

Combining Markov's inequality with the definition of subgaussian, we can show the following tail bound for subgaussian random variables.

Lemma 2.2 ([36], Proposition 2.6.6) *Let X be a centered random variable. X is subgaussian with parameter σ iff for any $c > 0$ it holds*

$$\Pr\left[|X| \geq c\right] \leq 2\exp\left(-\pi\frac{c^2}{\sigma^2}\right) \tag{2.4}$$

Subgaussian random variables behaves quite well when they are independent.

Lemma 2.3 ([36], Theorem 2.7.3) *Let $X_1 \leftarrow \mathcal{X}_{\sigma_1}, \ldots, X_n \leftarrow \mathcal{X}_{\sigma_n}$ be subgaussian random variables and $\mathbf{v} \in \mathbb{R}^n$. Consider $\mathbf{X}^\top = [X_1, \ldots, X_n]$ and $\sigma^\top = [\sigma_1, \ldots, \sigma_n]$. If $X_1, \ldots, X_n$ are independent, then $\langle \mathbf{X}, \mathbf{v} \rangle$ is subgaussian with parameter $\sqrt{\sum_{i=1}^n \left(\mathbf{v}[i]\right)^2 \sigma_i^2}$.*
In particular, when all subgaussians have the same parameter σ, we get that $\langle \mathbf{X}, \mathbf{v} \rangle$ is subgaussian with parameter $\|\mathbf{v}\|\sigma$.

It can also be shown that centered, t-bounded probability distributions are subgaussians with parameter $t\sqrt{2\pi}$.
We now extend the definition of subgaussian probability distribution to random vectors. We say that a probability distribution over $\mathbb{R}^n$ is subgaussian with parameter $\sigma > 0$ if $\mathbf{x} \leftarrow \mathcal{X}$ is centered, and for every $t \in \mathbb{R}$ and any unitary vector $\mathbf{u} \in \mathbb{R}^n$ it holds

$$\mathbb{E}\left[\exp\left(2\pi t\langle \mathbf{x}, \mathbf{u}\rangle\right)\right] \leq \exp\left(\pi\sigma^2 t^2\right). \tag{2.5}$$

In particular, by taking the inner product with any vector of the canonical basis of $\mathbb{R}^n$ we get that each coordinate of a subgaussian random vector is a subgaussian random variable with the same parameter. On the other hand, by using *Lemma 2.3* we can prove that if $X_1, \ldots, X_n \leftarrow \mathcal{X}_\sigma$ are independent, then the random vector $\mathbf{x} := [X_1, \ldots, X_n]^\top$ is a subgaussian vector with parameter σ. Similarly to what we have done for random variable over $\mathbb{R}$, we now state a tail bound for subgaussian random vector.

Lemma 2.4 (Lemma 2.2. [27]) *Let $X_1, \ldots, X_n \leftarrow \mathcal{X}_\sigma$ independent random variables. For any $c \geq 8n\sigma^2/\pi$ we have that*

$$\Pr\left[\sum_{i=1}^{n} X_i^2 > c\right] \leq \exp\left(-\pi\frac{c}{4\sigma^2}\right). \tag{2.6}$$

We have seen that the class of subgaussian probability distributions is quite broad and it satisfies several properties also known to be satisfied by the class of gaussian probability distributions. The latter class also fulfils another interesting properties: Let (X, Y) be jointly gaussian random variables. Then, for any y, the conditional distribution $X|Y = y$ is gaussian. The same unfortunately is not true in general for the class of subgaussian random variables.

The security reduction we want to generalize deals with a specific conditional distribution. Consider two independent random variables $X \leftarrow \mathcal{X}$ and $Y \leftarrow \mathcal{Y}$, and define the random variable $Z = X + Y$. We are now interested in analysing the random variable $X|Z = z$. Keeping gaussian distributions in mind, we formalize our assumption on the error distribution considered.

Assumption 2.5 *Let $\mathcal{X}$ and $\mathcal{Y}$ be two discrete subgaussian probability distributions with parameter σ_x and σ_y respectively. Let $X \leftarrow \mathcal{X}$ and $Y \leftarrow \mathcal{Y}$ be two independent samples, and let $Z := X + Y$. For every $z \in Supp(Z)$ such that $\Pr[Z = z]$ is non-negligible, there exists a probability distribution $\mathcal{W}_z$ such that*

1. *$\mathcal{W}_z$ is subgaussian with parameter $\sigma_w \leq C\sigma_x\sigma_y/\sqrt{\sigma_x^2 + \sigma_y^2}$, for some constant $C > 0$;*
2. *The distributions $X|Z = z$ and $\mathcal{W}_z + z\sigma_x^2/\left(\sigma_x^2 + \sigma_y^2\right)$ are within negligible statistical distance.*

This extra assumption is not artificial as it looks like. Indeed, in Prop. A.1 and Prop. A.2 we prove that centered binomials and discrete gaussians belong to this subclass. These examples cover most LPR-like schemes in the literature [2, 8, 29, 35], including ML-KEM which uses a centered binomial distribution.

2.4 Learning With Errors

We now recall the notions of lattice and ideal lattice together with the related hardness assumptions.

A lattice is a discrete additive subgroup of $\mathbb{R}^n$. We are interested only in full-rank lattices. This means that such a lattice can be described as

$$\Lambda = \mathcal{L}(\mathbf{B}) := \left\{ \sum_{i=1}^{n} z_i \mathbf{b}_i : z_i \in \mathbb{Z} \right\},$$

where $\mathbf{B} = [\mathbf{b}_1 | \dots | \mathbf{b}_n]$ is a basis of $\mathbb{R}^n$. Two bases $\mathbf{B}$ and $\mathbf{B}'$ generate the same lattice iff there exists a unimodular matrix $\mathbf{U}$ such that $\mathbf{B}' = \mathbf{UB}$.

A central hardness assumption based on lattices is the Learning With Errors (LWE) problem. It was first introduced by Regev in [33] and it is one of the most promising problems used to build post-quantum cryptographic protocols. The problem itself is quite simple to describe. Its decision version asks to distinguish between independent samples from the uniform distribution over $\mathbb{Z}_q^n \times \mathbb{Z}_q$ and the same number of independent samples from what is called LWE distribution. We briefly define the distribution and formalize the decision version of the problem.

Definition 2.6 (LWE distribution) *For a secret vector $\mathbf{s} \in \mathbb{Z}_q^n$ and a distribution $\mathcal{X}$ over $\mathbb{Z}_q$, a sample from the LWE distribution $A_{\mathbf{s},\mathcal{X}}$ is generated by choosing $\mathbf{a} \in \mathbb{Z}_q^n$, $e \leftarrow \mathcal{X}$, and outputting the pair $(\mathbf{a}, b = \langle \mathbf{a}, \mathbf{s} \rangle + e)$.*

Definition 2.7 (Decision LWE problem) *Given a modulus q and a probability distribution $\mathcal{X}$ over $\mathbb{Z}_q$, the decision version of the LWE problem (*DLWE*), denoted by $\mathsf{DLWE}_{q,\mathcal{X}}$, is to distinguish with non-negligible advantage between independent samples from $A_{\mathbf{s},\mathcal{X}}$ and the same number of uniformly random and independent samples from $\mathbb{Z}_q^n \times \mathbb{Z}_q$.*

Starting with [33], which provides a quantum reduction from the GapSVP and SIVP problems to the LWE problem, several follow-up works [10, 30] made the reduction classical. The main drawback of this hardness assumption is that protocols based on it were not so efficient. To address this inefficiency, several algebraic variants of LWE were introduced. Before defining them, we briefly recall the notion of ideal lattices.

Let $\mathbb{K}_m$ be the mth cyclotomic field and $\mathcal{R}$ its ring of integers. An ideal $\mathcal{I} \subset \mathcal{R}$ is an additive subgroup that is closed under multiplication by elements of $\mathcal{R}$, that is, $r \cdot a \in \mathcal{I}$ for any $a \in \mathcal{I}$ and $r \in \mathcal{R}$. An ideal $\mathcal{I} \subset \mathcal{R}$ always has a $\mathbb{Z}$-basis of cardinality n. In particular, if $\mathcal{I} = \langle a \rangle$ and $\mathbf{B}$ is a $\mathbb{Z}$-basis for $\mathcal{R}$, then $a \cdot \mathbf{B}$ is a $\mathbb{Z}$-basis for $\mathcal{I}$. We can extend the notion of ideals in $\mathbb{K}_m$ as follows: we say that $\mathcal{J} \subset \mathbb{K}_m$ is a fractional ideal if there exists $a \in \mathcal{R} \setminus \{0\}$ such that $a\mathcal{J} \subset \mathcal{R}$ is an ideal. As before, every fractional ideal $\mathcal{J}$ admits a $\mathbb{Z}$-basis of cardinality n. Given a fractional ideal $\mathcal{I} \subset \mathbb{K}_m$, we define an ideal lattice as the image of $\mathcal{I}$ through the coefficient embedding. Indeed by considering a $\mathbb{Z}$-basis $\{b_1, \dots, b_n\}$ of $\mathcal{I}$, its image through the embedding gives us the basis of a lattice

in $\mathbb{R}^n$. Ideal lattices provide a compact algebraic structure that naturally leads to a ring variant of LWE problem.

We now define the Ring-LWE (RLWE) problem and describe the worst-case hardness reduction first presented in [26]. Both the definition and the next statements are adaptations of the ones in [26, 27] due to the use of the different embedding. Moreover, since we are considering the case where m is a power of 2, we don't need to distinguish between $\mathcal{R}$ and its dual, since it is an integer scaling of $\mathcal{R}$. Given a modulus q, we denote $\mathcal{R}/q\mathcal{R}$ simply by $\mathcal{R}_q$.

Definition 2.8 (RLWE distribution) *For a secret $s \in \mathcal{R}_q$ and a distribution $\mathcal{X}$ over $\mathcal{R}$, a sample from the RLWE distribution $A_{s,\mathcal{X}}$ is generated by choosing $a \leftarrow_{\$} \mathcal{R}_q$, $e \leftarrow \mathcal{X}$, and outputting the pair $(a, b = as + e)$.*

Definition 2.9 (Decision RLWE problem) *Given a ring $\mathcal{R}$, a modulus q, and a probability distribution $\mathcal{X}$ over $\mathcal{R}_q$, the decision version of the RLWE problem (*DRLWE*), denoted $\mathcal{R}$-DRLWE$_{q,\mathcal{X}}$ is to distinguish with non-negligible advantage between independent samples from $A_{s,\mathcal{X}}$ and the same number of uniformly random and independent samples from $\mathcal{R}_q \times \mathcal{R}_q$.*

For simplicity, we will write DRLWE$_{q,\mathcal{X}}$ instead of $\mathcal{R}$-DRLWE$_{q,\mathcal{X}}$ when the ring we are using is clear from the context. The DRLWE problem reduces from worst-case problems over ideal lattices such as the Shortest Independent Vectors Problem (SIVP) and the Shortest Vector Problem (SVP) [26].

It can be shown that this version of DRLWE is at least as hard as its "normal" form [3, 25], that is the variant of the problem where the secret s is also sampled according to the probability distribution $\mathcal{X}$.
A further generalization of RLWE considers modules over the ring of integers, leading to the Module-LWE (MLWE) problem as has been shown in [25]. Let d be a positive integer, $\mathcal{R}$ be the ring of integers of $\mathbb{K}_m$, $n = \varphi(m)$.

Definition 2.10 (MLWE distribution) *For a secret vector $S \in \mathcal{R}_q^d$ and a distribution $\mathcal{X}$ over $\mathcal{R}$, a sample from the MLWE distribution $A_{S,\mathcal{X}}$ is generated by choosing $A \leftarrow_{\$} \mathcal{R}_q^d$, $e \leftarrow \mathcal{X}$, and outputting the pair $(A, b = \langle A, S \rangle + e)$.*

Definition 2.11 (Decision MLWE problem) *Given a module $\mathcal{R}^d$, a modulus q, and a probability distribution $\mathcal{X}$ over $\mathcal{R}_q$, the decision version of the MLWE problem (*DMLWE*), denoted DMLWE$_{q,\mathcal{X}}$ is to distinguish with non-negligible advantage between independent samples from $A_{S,\mathcal{X}}$ and the same number of uniformly random and independent samples from $\mathcal{R}_q^d \times \mathcal{R}_q$.*

As in the ring setting, the hardness of MLWE can be reduced from worst-case lattice problems such as the Module Generalized Independent Vectors Problem (M-GIVP) [25]. These results justify the use of RLWE and MLWE as efficient alternatives to standard LWE while retaining strong security guarantees.

$\underline{\mathsf{FFP\text{-}NG}_{\Pi}(\mathcal{A})}$:	$\underline{\mathsf{FCO}_{\beta}(\mu, r)}$:	#One query
01 $(sk_0, pk_0) \leftarrow_{\$} \mathsf{KeyGen}()$	06 $c \leftarrow \mathsf{Enc}(pk_\beta, \mu; r)$	
02 $(sk_1, pk_1) \leftarrow_{\$} \mathsf{KeyGen}()$	07 $\mu' := \mathsf{Dec}(sk_\beta, c)$	
03 $\beta \leftarrow_{\$} \{0,1\}$	08 **return** $[\![\mu = \mu']\!]$	
04 $\beta' \leftarrow \mathcal{A}^{\mathsf{FCO}_\beta}(pk_0)$		
05 **return** $[\![\beta = \beta']\!]$		

Fig. 2.1. The FFP-NG game against the PKE scheme $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$.

2.5 Find Failing Plaintext - Non Generic

In this subsection we recall the FFP-NG security notion, state some basic properties referring to [28] for the proofs, and describe how these properties will be turned into assumptions about FFP-NG adversaries.

In [20], the authors introduce a novel framework to analyse the impact of decryption failures on the security of PKE schemes. They introduce a family of security games called **Find Failing Plaintext (FFP)**. In this work we are interested in one member of this family, namely the Find Failing Plaintext - Non Generic (FFP-NG) notion. The FFP-NG game is described in Fig. 2.1.

In this game, a challenger honestly generates two independent key pairs (sk_0, pk_0) and (sk_1, pk_1) and provides the adversary $\mathcal{A}$ with the public key pk_0. The adversary gets a single query to a Failure-Checking Oracle (FCO) which uses one of the two key pairs, selected by sampling a bit β uniformly at random. This oracle takes as input a message-randomness pair (μ, r), and checks if it triggers a decryption failure with respect to the key pair (sk_β, pk_β). The goal of $\mathcal{A}$ is to understand which key pair has been used by FCO using the message-randomness pair (μ, r) chosen by the adversary and the response given by the oracle. In other words, the adversary should find a message-randomness pair that triggers a decryption failure with a non-negligible difference with respect to the given key pair compared to an independent one. By using the game described in Fig. 2.1, we define the advantage of the adversary $\mathcal{A}$ against a PKE scheme $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ as

$$\mathrm{Adv}_{\Pi}^{\mathsf{FFP\text{-}NG}}(\mathcal{A}) = \left|\Pr\left[\mathsf{FFP\text{-}NG}_{\Pi}(\mathcal{A}) = 1\right] - 1/2\right|.$$

The next result is a straightforward generalization of [28, Proposition 3].

Proposition 2.12 *Let Π be a PKE scheme and let $\mathcal{A}$ be an* FFP-NG *adversary against Π. If $\mathcal{A}$ chooses the message-randomness (μ, r) pair independent of the given key pk_0 then*

$$Adv_{\Pi}^{\mathsf{FFP\text{-}NG}}(\mathcal{A}) = 0.$$

The proof of the proposition is the same as [28, Proposition 3]. The key observation is that the two key pairs used in the FFP-NG game are generated honestly, thus they are identically distributed. If an adversary queries the failure-checking oracle using a message-randomness pair that is independent of the

given key, the probability of triggering a decryption failure is independent of the sampled bit β. Therefore, the adversary has no advantage in distinguishing which key pair has been used by the failure-checking oracle. This proposition will be used in the reduction to better analyse the error distribution appearing in the decryption failure conditions derived later.

Furthermore, if the probability of triggering a decryption failure is in general quite low, the adversary has limited options to gain a significant advantage. Indeed, when dealing with PKE schemes that are correct with overwhelming probability, the only strategy the adversary has to gain a significant advantage is to output a message-randomness pair that triggers a decryption failure with higher probability with respect to the given key than with respect to an independent one.

3 Security Reductions from Ring LWE

This section is structured in the following way:

1. In Sect. 3.1, we define the PKE scheme introduced in [26], state its correctness and IND-CPA security, and describe in detail the related decryption failure condition.

2. In Sect. 3.2, we prove our reduction from DRLWE to FFP-NG using the LPR scheme as underlying PKE scheme. The error distributions are taken from the subclass of subgaussian distributions satisfying Assumption 2.5. This allows us to get a reduction for both discrete gaussians and centered binomials within a unified framework.

Throughout the section, we will use the following set of parameters. Let $m > 4$ be a power of 2, $n = \varphi(m) = m/2$, $\Phi_m(X)$ be the m-th cyclotomic polynomial, $\mathbb{K}_m$ be the m-th cyclotomic field, and $\mathcal{R}$ its ring of integers. Let q be an odd prime such that $q = 1 \mod m$, and let t be an integer such that $1 < t < \sqrt{\log n}$. The value n is the security parameter.

3.1 The LPR public key encryption scheme

In [26], the authors describe a PKE scheme based on the hardness of the RLWE problem. We refer to it as the LPR.PKE scheme. This protocol is the extension of the Dual Regev scheme [17] to the ring setting and it provides the basis for several constructions [2, 8, 29, 35]. This is the reason why we decided to use it as the underlying PKE scheme. To keep our description as general as possible, we avoid picking a specific probability distribution. Instead, we consider classes of probability distributions.

We define LPR.PKE = (LPR.KeyGen, LPR.Enc, LPR.Dec) in Fig. 3.1, where $\mathcal{R}_2$ is the message space, $\mathcal{X}_\sigma$ is a subgaussian probability distribution over $\mathcal{R}_q$, and $\mathcal{X}$ is a t-bounded probability distribution over $\mathcal{R}_q$.

LPR.KeyGen():	LPR.Enc (pk, μ):	LPR.Dec(sk, c):
01 $a \leftarrow_{\$} \mathcal{R}_q$	06 $r, e_1, e_2 \leftarrow \mathcal{X}$	10 $(c_1, c_2) = c$
02 $s, e \leftarrow \mathcal{X}_\sigma$	07 $c_1 := ar + e_1 \mod q$	11 $\mu' := \lfloor c_2 - sc_1 \mod q \rceil_2$
03 $b := as + e$	08 $c_2 := br + e_2 + \lfloor \mu \rceil_q \mod q$	12 **return** μ'
04 $sk := s$, $pk := (a, b)$	09 **return** $c := (c_1, c_2)$	
05 **return** (pk, sk)		

Fig. 3.1. The PKE scheme LPR.PKE introduced in [26], where $\lfloor \cdot \rceil_q = \lfloor \cdot \rceil_{2 \to q}$ and $\lfloor \cdot \rceil_2 = \lfloor \cdot \rceil_{q \to 2}$ are applied coordinate-wise.

Let $\mu \in \mathcal{R}_2$ be a message. To obtain a decryption failure in Π, the following equation should hold.

$$\left\lfloor er - se_1 + e_2 + \lfloor \mu \rceil_q \right\rceil_2 \neq \mu.$$

Let us consider $z = er - se_1 + e_2$ the total noise, and let $\mathbf{z}$ be its coefficient vector. A decryption failure means that one of the coefficients of $er - se_1 + e_2 + \lfloor \mu \rceil_q$ gets rounded to the wrong value. In other words, there is an index $k \in \{0, \ldots, n-1\}$ such that

$$\left| \frac{2}{q} \left(z[k] + \left\lfloor \frac{q}{2} \mu[k] \right\rceil \right) - \mu[k] \right| \geq \frac{1}{2}.$$

By using the triangle inequality, we get

$$\frac{2}{q} |z[k]| + \left| \frac{2}{q} \left\lfloor \frac{q}{2} \mu[k] \right\rceil - \mu[k] \right| \geq \frac{1}{2}. \tag{3.1}$$

By definition of rounding, we also get

$$\left| \frac{2}{q} \left\lfloor \frac{q}{2} \mu[k] \right\rceil - \mu[k] \right| = \frac{2}{q} \left| \left\lfloor \frac{q}{2} \mu[k] \right\rceil - \frac{q}{2} \mu[k] \right| \leq \frac{1}{q}. \tag{3.2}$$

Combining Eq. (3.1) and Eq. (3.2) we have $|z[k]| \geq (q-2)/4$. This means that if a failure occurs, then $\|z\|_\infty \geq (q-2)/4$. We obtained a necessary condition for decryption failure. It would be helpful to also have a sufficient condition for decryption failure. Using the reverse triangle inequality we can show that if $\|z\|_\infty \geq (q+2)/4$ then we get a decryption failure. The reasoning above provides a proof of the following result.

Lemma 3.1 (Decryption Failure Conditions) *Let $z = er - se_1 + e_2$ be the total error obtained during decryption. We have*

1. *If a decryption failure occurs $\implies$ $\|z\|_\infty \geq (q-2)/4$.*
2. *If $\|z\|_\infty \geq (q+2)/4 \implies$ a decryption failure occurs.*

We want to show that the scheme is correct with overwhelming probability. To do so, we will exploit some of the results about subgaussian probability distributions described in Sect. 2.

Lemma 3.2 (Correctness) *Let $\mathcal{X}_\sigma$ be a subgaussian distribution, and let $\mathcal{X}$ be a t-bounded probability distribution over $\mathcal{R}_q$. If $q \in \Omega(n\log(n))$ and $\sigma \in \Theta\left(\sqrt{\log(n)}\right)$, then the PKE scheme described in Fig. 3.1 is correct with overwhelming probability.*

Proof. According to Lemma 3.1, if the ℓ_∞ norm of the total noise is smaller than $(q-2)/4$ then the decryption is correct. Instead of proving

$$\Pr\left[\|re - se_1 + e_2\|_\infty < \frac{q-2}{4}\right] = 1 - negl(n),$$

we will prove that

$$\Pr\left[\|re - se_1 + e_2\|_\infty \geq \frac{q-2}{4}\right] \in negl(n),$$

where, in both cases, the probability is taken over the randomness of $\mathsf{LPR.KeyGen}$ and $\mathsf{LPR.Enc}$.

We define $\mathbf{v} := \left[\mathbf{e}^\top, -\mathbf{s}^\top\right]^\top$ and $\mathbf{w}_{(k)} := \left[{\mathbf{r}_{(k)}}^\top, {\mathbf{e_1}_{(k)}}^\top\right]^\top$, where $\mathbf{e}$, $\mathbf{s}$, $\mathbf{r}$, and $\mathbf{e_1}$ are the coefficient vectors of e, s, r, and e_1 respectively. We can write

$$\Pr\left[\|re - se_1 + e_2\|_\infty \geq \frac{q-2}{4}\right] = \Pr\left[\exists k : |\langle \mathbf{v}, \mathbf{w}_{(k)}\rangle + \mathbf{e_2}[k]| \geq \frac{q-2}{4}\right].$$

Let us fix an index $k \in \{0, \ldots, n-1\}$. By using Lemma 2.3, we have that $\langle \mathbf{v}, \mathbf{w}_{(k)}\rangle$ is subgaussian with parameter $\sigma t\sqrt{2n}$. So, the tail bound in Lemma 2.2 gives us

$$\begin{aligned}\Pr\left[|\langle \mathbf{v}, \mathbf{w}_{(k)}\rangle + \mathbf{e_2}[k]| \geq \frac{q-2}{4}\right] &\leq \Pr\left[|\langle \mathbf{v}, \mathbf{w}_{(k)}\rangle| \geq \frac{q}{4} - \frac{1}{2} - t\right] \\ &\leq 2\exp\left(-\pi\left(\frac{q-4t-2}{4\sqrt{2n}\sigma t}\right)^2\right).\end{aligned}$$

To prove that the left-hand side of the inequality is negligible in n, we can prove that

$$\frac{q-4t-2}{4t\sigma\sqrt{2n}} \in \omega\left(\sqrt{\log(n)}\right).$$

This follows by the fact that

$$\frac{4t+2}{4t\sigma\sqrt{2n}} \in \Theta\left(\frac{1}{\sqrt{n\log(n)}}\right),$$

and

$$\frac{q}{4t\sigma\sqrt{2n}} \in \omega\left(\sqrt{\log(n)}\right),$$

where we use the assumption on q, t, and σ.

The obtained bound is independent of the index k; hence by applying the union bound, we get

$$\Pr\left[\|re - se_1 + e_2\|_\infty > \frac{q-2}{4}\right] \in negl\,(n)\,.$$

□

Lemma 3.3 (IND-CPA security) *The PKE scheme described in Fig. 3.1 is* IND-CPA *secure assuming the hardness of* $\mathsf{DRLWE}_{q,\mathcal{X}_\sigma}$ *and* $\mathsf{DRLWE}_{q,\mathcal{X}}$.

The proof follows from two applications of the RLWE hardness assumption in its normal form, [27, Lemma 2.24]. The first application is used to prove that the public key is indistinguishable from uniform. The second application is used to prove that the adversary's point of view is indistinguishable from uniform.

3.2 Security reduction with Subgaussian distributions

We now turn our attention to the security reduction from DRLWE to FFP-NG. The idea is to adapt the reduction presented in [28] while generalizing the family of probability distributions allowed.

Once we identify the coefficient embedding as the appropriate one to use in this setting and properly define the decryption failure condition, extending the analysis from plain LWE to RLWE becomes relatively simple. Our main goal, however, is to make the result as general as possible, which requires selecting an underlying PKE scheme that is widely used in practice and supports a broad class of error distributions.

In the previous subsection, we considered subgaussian distributions as key generation error distributions. This family turns out to be too broad for this adaptation. Indeed, we need to analyse conditional distributions, and we have already observed that subgaussianity is not preserved in general under conditioning. For this reason, we introduced in Assumption 2.5 a property that allows us to control the relevant conditional probabilities.

To simplify the statements in this section, we collect all assumptions on the main parameters and error distributions in Fig. 3.2.

We are now ready to state the main theorem of this section.

Theorem 3.4 (from DRLWE to FFP-NG). *Let n, q, t, $\mathcal{X}$, $\mathcal{X}_\sigma$, $\mathcal{X}_{\sigma'}$, and $\mathcal{X}_{\hat{\sigma}}$ satisfy the assumptions in Fig. 3.2. For all* FFP-NG *adversaries $\mathcal{A}$ against the* LPR.PKE *scheme, with errors distribution $\mathcal{X}_\sigma$ and randomness $\mathcal{X}$, denoted by Π, there exists an adversary $\mathcal{B}$ that solves the* DRLWE *problem with error distribution $\mathcal{X}_{\hat{\sigma}}$ such that*

$$Adv_{\Pi}^{\mathsf{FFP\text{-}NG}}(\mathcal{A}) \leq Adv_{\Pi}^{\mathsf{DRLWE}}(\mathcal{B}) + \Gamma(n)\,, \tag{3.3}$$

where Γ is negligible in n.

Assumptions: let C_q, c_σ, C_σ some positive constant values, let $1 < t < \sqrt{\log(n)}$ an integer, let $q = q(n)$ a prime number, and let $\mathcal{X}$, $\mathcal{X}_\sigma$, $\mathcal{X}_{\sigma'}$, and $\mathcal{X}_{\hat{\sigma}}$ be subgaussian distributions, for which $\exists n_0 > 0$ such that for all $n > n_0$ it holds

- $q \geq C_q n \log(n)$ and $q = 1 \mod 2n$;
- $\mathcal{X}$ a t-bounded probability distribution;
- if $x' \leftarrow \mathcal{X}_{\sigma'}$ and $\hat{x} \leftarrow \mathcal{X}_{\hat{\sigma}}$ are independent samples, then $(x' + \hat{x}) \sim \mathcal{X}_\sigma$;
- $\sigma^2 = \sigma'^2 + \hat{\sigma}^2$, and $\sigma' < \hat{\sigma}$;
- $c_\sigma \sqrt{\log(n)} \leq \sigma$, σ', $\hat{\sigma} \leq C_\sigma \sqrt{\log(n)}$;
- $\sqrt{\pi} C_q > 32 C_\sigma$.
- $\mathcal{X}_{\sigma'}$ and $\mathcal{X}_{\hat{\sigma}}$ satisfy Assumption 2.5

Fig. 3.2. Parameters assumptions for Thm. 3.4.

The function Γ captures the loss of the reduction. Its exact expression is cumbersome and arises from lengthy computations. Therefore, we present only its asymptotic behaviour here and refer to the supplementary material for the concrete bound.

Before getting into the details of the proof, let us briefly describe the idea behind it. The reduction exploits the FFP-NG adversary as follows. Given an instance $pk = (a, b)$ of the DRLWE problem, the reduction samples a pair of small errors s', e', and uses them to tweak the given sample. The modification is such that if the sample was uniform, it remains uniform, whereas if it was an RLWE sample, it remains an RLWE sample but with a slightly modified error distribution. The reduction then calls the FFP-NG adversary on input the modified sample pk', obtains a message-randomness pair, and uses the randomness and the known errors to set a threshold θ. It is important to stress that the DRLWE adversary simulates the failure-checking oracle until the FFP-NG adversary queries it.

The intuition is that when the FFP-NG adversary receives a uniform pair, the randomness (r, e_1, e_2) used to query the failure-checking oracle is independent of the added error (s', e'). In high dimension, these values have a tendency to be orthogonal, and the norm of their inner product will not exceed the threshold. We prove in Lemma 3.5.

On the other hand, when the adversary receives an LWE sample and succeeds in causing a decryption failure, the randomness (r, e_1, e_2) is aligned with the error, and thus also with the added error $\mathbf{e}'$. This means that their inner product is likely to exceed the threshold. We prove this in Lemma 3.6.

Proof. If $Adv_{\Pi}^{\mathsf{FFP\text{-}NG}}(\mathcal{A})$ is negligible in n, Eq. (3.3) holds trivially. In the following, we assume that $\mathcal{A}$ has non-negligible advantage. The DRLWE adversary $\mathcal{B}$ is defined in Fig. 3.3. Notice that when $\mathcal{B}$ gets a uniformly random pair (a, b), the modified public key $pk' = (a, b + b')$ is also uniformly random. In particular, the uniform b prevents $\mathcal{A}$ from obtaining any information about (s', e'). On the other

$\underline{\mathcal{B}(a,b)}$:
$s', e' \leftarrow \mathcal{X}_{\sigma'}$
$b' := as' + e'$
$pk' := (a, b + b')$
$(\mu, (r, e_1, e_2)) \leftarrow \mathcal{A}(pk')$
$\mathbf{w}^\top := \left[\mathbf{r}^\top, \mathbf{e_1}^\top\right]$, $\mathbf{v'}^\top := \left[\mathbf{e'}^\top, -\mathbf{s'}^\top\right]$, ${\mathbf{w}_{(k)}}^\top := \left[{\mathbf{r}_{(k)}}^\top, {\mathbf{e_1}_{(k)}}^\top\right]$
$\theta := (q\|\mathbf{w}\|) / (8t\sqrt{n})$
if $\exists k$ s.t. $|\langle \mathbf{v'}, \mathbf{w}_{(k)} \rangle + \mathbf{e_2}[k]| \geq (q - 4\theta - 2)/4$
  return 1 #RLWE
else
  return 0 #Uniform

Fig. 3.3. Adversary against the $\mathsf{DRLWE}_{q,\mathcal{X}_{\hat{\sigma}}}$ problem.

hand, when the pair (a, b) is a RLWE sample, we have $b+b' = a(s + s')+(e + e')$, which is still an RLWE sample.

To prove $\mathcal{B}$ has a non-negligible advantage in solving the DRLWE problem, we need to show that the following advantage is non-negligible in n:

$$\mathrm{Adv}_{\Pi}^{\mathsf{DRLWE}}(\mathcal{B}) := \left| \Pr_{(a,b)\leftarrow\mathsf{RLWE}} [\mathcal{B}(a,b) = 1] - \Pr_{(a,b)\leftarrow\$\mathcal{R}_q^2} [\mathcal{B}(a,b) = 1] \right|.$$

Given the tuple (r, e_1, e_2), for $k \in \{0, \ldots, n-1\}$ we define the event

$$\mathrm{FAIL}_k(r, e_1, e_2) := \left\{ |\langle \mathbf{v'}, \mathbf{w}_{(k)} \rangle + \mathbf{e_2}[k]| \geq \frac{q - 4\theta - 2}{4} \right\}, \tag{3.4}$$

where $\mathbf{w}_{(k)}$ and $\mathbf{v'}$ are defined in Fig. 3.3. By definition of $\mathcal{B}$, it suffices to prove that the following absolute difference is non-negligible in n:

$$\left| \Pr_{(a,b)\leftarrow\mathsf{RLWE}} [\exists k\colon \mathrm{FAIL}_k(r, e_1, e_2)] - \Pr_{(a,b)\leftarrow\$\mathcal{R}_q^2} [\exists k\colon \mathrm{FAIL}_k(r, e_1, e_2)] \right|. \tag{3.5}$$

We analyse the two probabilities separately: in Lemma 3.5 we prove that the second term in Eq. (3.5) is negligible in n, and in Lemma 3.6 we prove that the first term in Eq. (3.5) is non-negligible in n. Combining the two results completes the proof.

□

We start with bounding the second term in Eq. (3.5), i.e., we show that when the public key is uniformly random, it is hard for the adversary to consistently output randomness aligned with the added secret.

Lemma 3.5 (FAIL with uniformly random key) *Let $\mathbf{v'} \leftarrow \mathcal{X}_{\sigma'}$ be an error vector, and $(r, e_1, e_2) \leftarrow \mathcal{A}(pk')$ be the adversarially chosen encryption randomness. If (a, b) is sampled uniformly at random, then*

$$\Pr_{(a,b)\leftarrow\$\mathcal{R}_q^2} [\exists k\colon \mathrm{FAIL}_k(r, e_1, e_2)] \in negl(n). \tag{3.6}$$

Proof. Throughout the proof, we omit the subscripts of the probabilities, as they are clear from the statement. Since (a, b) is sampled uniformly at random, $pk' = (a, b + b')$ is also uniformly random. This means that the tuples (s', e') and (r, e_1, e_2) are independent. Thus, we can write

$$\Pr\left[\exists k\colon \mathrm{FAIL}_k\left(r, e_1, e_2\right)\right] \leq n \cdot \max_k \left\{\Pr\left[\ \mathrm{FAIL}_k\left(r, e_1, e_2\right)\right]\right\}.$$

We fix an index k and bound $\Pr\left[\ \mathrm{FAIL}_k\left(r, e_1, e_2\right)\right]$.

Recall the definition of $\mathbf{v}'$ and $\mathbf{w}_{(k)}$ in Fig. 3.3. As they are independent, the inner product $\langle \mathbf{v}', \mathbf{w}_{(k)} \rangle$ is subgaussian with parameter $\|\mathbf{w}\|\sigma'$. We use Lemma 2.2 with $c = \left(q - 4\theta - 4t - 2\right)/4$ and obtain

$$\Pr\left[\ \mathrm{FAIL}_k\left(r, e_1, e_2\right)\right] \leq \Pr\left[\left|\langle \mathbf{v}', \mathbf{w}_{(k)} \rangle\right| \geq c\right] \leq 2\exp\left(-\pi\left(\frac{q - 4\theta - 4t - 2}{4\|\mathbf{w}\|\sigma'}\right)^2\right).$$

We have obtained a bound that is independent of k.

Using the definitions of the parameters and θ, we write

$$\frac{q - 4\theta - 4t - 2}{4\|\mathbf{w}\|\sigma'} \geq \frac{q - 4t - 2}{4t\sigma'\sqrt{2n}} - \frac{\theta}{\|\mathbf{w}\|\sigma'} = \frac{q - 4t - 2}{4t\sigma'\sqrt{2n}} - \frac{q}{8t\sigma'\sqrt{n}} = \frac{\left(1 - \frac{1}{\sqrt{2}}\right)q - 4t - 2}{4t\sigma'\sqrt{2n}}$$

Similarly to what we have done in Lemma 3.2, we can prove that

$$\frac{q - 4\theta - 4t - 2}{4\|\mathbf{w}\|\sigma'} \in \omega\left(\sqrt{\log n}\right).$$

From this we get the final bound

$$\Pr\left[\exists k\colon \mathrm{FAIL}_k\left(r, e_1, e_2\right)\right] \leq n \cdot \exp\left(-\omega\left(\log(n)\right)\right)$$

which is negligible in n.

□

Now we bound the first term of Eq. (3.5). In particular, we show that, assuming the FFP-NG adversary has non-negligible advantage, the DRLWE adversary can exploit this adversary and the known error to gain information about the public key it received.

Lemma 3.6 (FAIL with RLWE key) *Let $\mathbf{v}' \leftarrow \mathcal{X}_{\sigma'}$ be an error vector, and $(r, e_1, e_2) \leftarrow \mathcal{A}\left(pk'\right)$ be the adversarially chosen encryption randomness. If (a, b) is sampled according to the RLWE distribution, then*

$$\Pr_{(a,b)\leftarrow \mathsf{RLWE}}\left[\exists k\colon \mathrm{FAIL}_k\left(r, e_1, e_2\right)\right] \notin negl\left(n\right). \tag{3.7}$$

This time the analysis is more involved than before. The public key (a, b) is sampled according to the RLWE distribution. This means that $a \leftarrow_{\$} \mathcal{R}_q$ and $b = a\hat{s} + \hat{e}$, with $\hat{s}$, $\hat{e}$ sampled according to $\mathcal{X}_{\hat{\sigma}}$. In this case, the modified public key pk' is equal to $(a, as + e)$, with $e = e' + \hat{e}$ and $s = s' + \hat{s}$ distributed according to $\mathcal{X}_{\sigma}$. The randomness (r, e_1, e_2) is not independent of the pair (e', s'). During the proof, we consider the following coefficient vectors

$$\mathbf{v}^{\top} := \left[\mathbf{e}^{\top}, -\mathbf{s}^{\top}\right], \ \mathbf{v'}^{\top} := \left[\mathbf{e'}^{\top}, -\mathbf{s'}^{\top}\right], \ \hat{\mathbf{v}}^{\top} := \left[\hat{\mathbf{e}}^{\top}, -\hat{\mathbf{s}}^{\top}\right],$$
$$\mathbf{w}^{\top} := \left[\mathbf{r}^{\top}, \mathbf{e_1}^{\top}\right], \ \mathbf{w}_{(k)}^{\top} := \left[\left(\mathbf{r}_{(k)}\right)^{\top}, \left(\mathbf{e_1}_{(k)}\right)^{\top}\right].$$

The proof works as follows. We use the observations made in Sect. 2.5 to narrow down the strategies available to the FFP-NG adversary $\mathcal{A}$ for achieving non-negligible advantage. In particular, $\mathcal{A}$ must trigger a decryption failure using the randomness it chooses. According to Lemma 3.1, conditioning on this event gives us information about the inner product between $\mathbf{w}_{(k)}$ and $\mathbf{v}$.

Combining these observations, we find ourselves in a situation where it suffices to show that the inner product between $\mathbf{w}_{(k)}$ and $\hat{\mathbf{v}}$ is large only with negligible probability. To obtain independence between these two random vectors, we condition on the value taken by $\mathbf{v}$. Thanks to Assumption 2.5, we can still characterize the resulting conditional distribution.

The final step is to introduce a second threshold γ. It is defined so that $|\mathbf{v}|$ exceeds γ only with negligible probability. At the same time, when $|\mathbf{v}| < \gamma$, the inner product between $\mathbf{w}_{(k)}$ and $\hat{\mathbf{v}}$ is large only with negligible probability. This is exactly what we want to prove.

Proof. Throughout the proof, we omit the subscripts of the probabilities, as they are clear from the statement. We need to characterize $\mathcal{A}$'s winning condition and express it in terms of independent tuples to facilitate the analysis. Let us define the event

$$\mathrm{WIN}\,(\mathcal{A}) = \{\mathcal{A} \text{ gets the same } pk \text{ used by } \mathsf{FCO} \text{ and wins the } \mathsf{FFP\text{-}NG} \text{ game}\},$$

and recall the definition of $\mathrm{FAIL}_k\,(r, e_1, e_2)$ given in Eq. (3.4). Using these events we can write

$$\Pr\left[\exists k\colon \mathrm{FAIL}_k\,(r, e_1, e_2)\right] \geq \Pr\left[\exists k\colon \mathrm{FAIL}_k\,(r, e_1, e_2)\,|\mathrm{WIN}\,(\mathcal{A})\right] \cdot \Pr\left[\mathrm{WIN}\,(\mathcal{A})\right].$$

Exploiting again Prop. 2.12 and the discussion at the end of Sect. 2.5, the event $\mathrm{WIN}\,(\mathcal{A})$ implies that $\mathcal{A}$ triggers a decryption failure for the given public key using the randomness (r, e_1, e_2). In turn, this means

$$\exists\ k : |\langle \mathbf{v}, \mathbf{w}_{(k)}\rangle + \mathbf{e_2}\,[k]\,| \geq \frac{q-2}{4} \implies \max_k\{|\langle \mathbf{v}, \mathbf{w}_{(k)}\rangle + \mathbf{e_2}\,[k]\,|\} \geq \frac{q-2}{4}.$$

Using the definition of $\mathbf{e}$ and $\mathbf{s}$, and the triangle inequality, we get

$$\max_k\{|\langle \mathbf{v}, \mathbf{w}_{(k)}\rangle + \mathbf{e_2}\,[k]\,|\} \leq \max_k\{|\langle \mathbf{v'}, \mathbf{w}_{(k)}\rangle + \mathbf{e_2}\,[k]\,|\} + \max_k\{|\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle|\}.$$

Thanks to this inequality, we obtain

$$\Pr\left[\exists k\text{: FAIL}_k\left(r, e_1, e_2\right) | \text{WIN}\left(\mathcal{A}\right)\right] \geq \Pr\left[\max_k\{|\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle|\} < \theta \;\middle|\; \text{WIN}\left(\mathcal{A}\right)\right].$$

To prove the lemma, it suffices to prove that this probability is non-negligible in n. To do so, we can prove

$$\Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta \;\middle|\; \text{WIN}\left(\mathcal{A}\right)\right] \in negl\left(n\right).$$

Exploiting that $\mathcal{A}$ has non-negligible advantage in the FFP-NG game, it is enough to prove

$$\Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta\right] \in negl\left(n\right). \tag{3.8}$$

Observe that, by conditioning on $\mathbf{v} = \mathbf{v}^* \in \mathbb{R}^n$, we get that $\hat{\mathbf{v}}$ and $\mathbf{w}_{(k)}$ are independent for every k. Furthermore, thanks to Assumption 2.5 there is a random variable $\overline{\mathbf{v}} \leftarrow \mathcal{X}_{\overline{\sigma}}$ such that $\hat{\mathbf{v}}|\mathbf{v} = \mathbf{v}^*$ and $\overline{\mathbf{v}} + \mathbf{v}^*\hat{\sigma}^2/\sigma^2$ are within negligible statistical distance. We want to distinguish among good conditioning and bad one. Here, good means that the vector $\mathbf{v}^*$ is not too far from the origin, while bad means that it has high ℓ_2 norm. So, we would like to define a threshold γ such that

1. $\Pr\left[\|\mathbf{v}\| \geq \gamma\right] \in negl\left(n\right)$;
2. if $\|\mathbf{v}^*\| \leq \gamma \implies \Pr\left[|\langle \overline{\mathbf{v}} + \mathbf{v}^*\hat{\sigma}^2/\sigma^2, \mathbf{w}_{(k)}\rangle| \geq \theta \;\middle|\; \mathbf{v} = \mathbf{v}^*\right] \in negl\left(n\right)$.

We start by setting the condition to fulfil $\Pr\left[\|\mathbf{v}\| \geq \gamma\right] \in negl\left(n\right)$. Thanks to Lemma 2.4 we know that

$$\Pr\left[\|\mathbf{v}\| > \gamma\right] \leq \exp\left(-\pi\frac{\gamma^2}{4\sigma^2}\right) \tag{3.9}$$

as soon as $\gamma\sqrt{\pi} \geq 2\sigma\sqrt{2n}$. Thus, by setting $\gamma := 4\sigma\sqrt{n/\pi}$, the right-hand side of Eq. (3.9) is negligible in n.

Thanks to this threshold we can write

$$\begin{aligned}\Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta\right] &= \Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta \;\middle|\; \|\mathbf{v}\| \geq \gamma\right]\Pr\left[\|\mathbf{v}\| \geq \gamma\right] \\ &\quad + \Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta \;\middle|\; \|\mathbf{v}\| < \gamma\right]\Pr\left[\|\mathbf{v}\| < \gamma\right] \\ &\leq \Pr\left[\|\mathbf{v}\| \geq \gamma\right] + \Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta \;\middle|\; \|\mathbf{v}\| < \gamma\right].\end{aligned}$$

We have already bounded the first term. We need to check if the threshold allows us to bound also the second one. Here we use Assumption 2.5 as follows

$$\begin{aligned}&\Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta \;\middle|\; \|\mathbf{v}\| < \gamma\right] \\ &\leq \sum_{\mathbf{v}^* \in B(0,\gamma)} \Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta \;\middle|\; \mathbf{v} = \mathbf{v}^*\right]\Pr\left[\mathbf{v} = \mathbf{v}^*\right] \\ &\leq \varepsilon + \sum_{\mathbf{v}^* \in B(0,\gamma)} \Pr\left[\exists\ k : |\langle \overline{\mathbf{v}} + \mathbf{v}^*\frac{\hat{\sigma}^2}{\sigma^2}, \mathbf{w}_{(k)}\rangle| \geq \theta\right]\Pr\left[\mathbf{v} = \mathbf{v}^*\right]\end{aligned}$$

where $B(0,\gamma)$ is the set of vectors with ℓ_2 norm smaller than γ, and $\varepsilon \in negl\,(n)$. If we could get a bound

$$\Pr\left[\exists\ k : |\langle \overline{\mathbf{v}} + \mathbf{v}^* \frac{\hat{\sigma}^2}{\sigma^2}, \mathbf{w}_{(k)}\rangle| \geq \theta\right] \in negl\,(n)$$

that is also independent of k and $\mathbf{v}^*$, we could bound $\sum \Pr\left[\mathbf{v} = \mathbf{v}^*\right] \leq 1$ and get

$$\Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta \ \middle|\ \|\mathbf{v}\| < \gamma\right] \in negl\,(n) \tag{3.10}$$

Recall that $\overline{\mathbf{v}}$ is subgaussian with parameter $\overline{\sigma} := (\hat{\sigma}\sigma')/\sigma$ and it is independent of $\mathbf{w}_{(k)}$, so the inner product is subgaussian with parameter $\|\mathbf{w}\|\overline{\sigma}$. Assume that $\mathbf{v}^* \in B(0,\gamma)$, we apply the triangle inequality and Lemma 2.2

$$\Pr\left[\left|\langle \overline{\mathbf{v}} + \mathbf{v}^* \frac{\hat{\sigma}^2}{\sigma^2}, \mathbf{w}_{(k)}\rangle\right| \geq \theta\right] \leq 2\exp\left(-\pi\left(\frac{\theta\sigma^2 - |\langle \mathbf{v}^*, \mathbf{w}_{(k)}\rangle|\hat{\sigma}^2}{\|\mathbf{w}\|\sigma^2\overline{\sigma}}\right)^2\right).$$

To apply the bound we must have $\theta\sigma^2 > |\langle \mathbf{v}^*, \mathbf{w}_{(k)}\rangle|\hat{\sigma}^2$ and to be useful we need

$$\frac{\theta\sigma^2 - |\langle \mathbf{v}^*, \mathbf{w}_{(k)}\rangle|\hat{\sigma}^2}{\|\mathbf{w}\|\sigma^2\overline{\sigma}} \in \omega\left(\sqrt{\log(n)}\right).$$

We apply the Cauchy-Schwarz inequality and get

$$\theta\sigma^2 - |\langle \mathbf{v}^*, \mathbf{w}_{(k)}\rangle|\hat{\sigma}^2 > \theta\sigma^2 - \gamma\|\mathbf{w}\|\hat{\sigma}^2 \geq \|\mathbf{w}\|\sigma^2\sqrt{n\log(n)}\left(\frac{C_q}{8} - \frac{4C_\sigma}{\sqrt{\pi}}\right),$$

which is greater than zero as soon as $\sqrt{\pi}C_q > 32C_\sigma$.

Since $\overline{\sigma} \in \Theta\left(\sqrt{\log(n)}\right)$, we have

$$\frac{\theta\sigma^2 - |\langle \mathbf{v}^*, \mathbf{w}_{(k)}\rangle|\hat{\sigma}^2}{\|\mathbf{w}\|\sigma^2\overline{\sigma}} \geq \frac{\sqrt{n\log(n)}}{\overline{\sigma}}\left(\frac{C_q}{8} - \frac{4C_\sigma}{\sqrt{\pi}}\right) \in \omega\left(\sqrt{\log(n)}\right),$$

Thanks to this, we have

$$\Pr\left[\exists\ k : |\langle \overline{\mathbf{v}} + \mathbf{v}^* \frac{\hat{\sigma}^2}{\sigma^2}, \mathbf{w}_{(k)}\rangle| \geq \theta\right] \leq n \cdot \exp\left(-\omega\left(\log(n)\right)\right)$$

which is negligible in n, and independent of k and $\mathbf{v}^*$. To conclude the proof, we combine Eq. (3.10) and Eq. (3.9), and prove that

$$\Pr\left[\exists\ k : |\langle \hat{\mathbf{v}}, \mathbf{w}_{(k)}\rangle| \geq \theta\right] \in negl\,(n).$$

□

KYBER.KeyGen():	KYBER.Enc (pk, μ):	KYBER.Dec(sk, c):
01 $\mathbf{A} \leftarrow_{\$} \mathcal{R}_q^{d \times d}$	07 $R, E_1 \leftarrow \mathcal{X}^d$	12 $(C_1, c_2) = c$
02 $S, E \leftarrow \mathcal{X}_\sigma^d$	08 $e_2 \leftarrow \mathcal{X}$	13 $\mu' := \lfloor c_2 - \langle S, C_1 \rangle \mod q \rceil_2$
03 $B := \mathbf{A}S + E$	09 $C_1^\top := R^\top \mathbf{A} + E_1^\top \mod q$	14 **return** μ'
04 $sk := S$	10 $c_2 := \langle B, R \rangle + e_2 + \lfloor \mu \rceil_q \mod q$	
05 $pk := (\mathbf{A}, B)$	11 **return** $c := (C_1, c_2)$	
06 **return** (pk, sk)		

Fig. 4.1. The functions $\lfloor \cdot \rceil_q : \mathcal{R}_2 \to \mathcal{R}_q$ and $\lfloor \cdot \rceil_2 : \mathcal{R}_q \to \mathcal{R}_2$ are the coordinate-wise rounding to the closest integer.

4 Security Reductions from Module LWE

This section is structured as follows:

1. In Sect. 4.1, we generalize the public key encryption scheme analysed in Sect. 3.1 to modules over cyclotomic rings. This scheme is a simplified version of the PKE scheme on which ML-KEM is built. We state its correctness and IND-CPA security, and describe the related decryption failure condition.

2. In Sect. 4.2, we show that the reduction presented in Sect. 3.2 can be extended with minimal changes assuming the hardness of DMLWE. In this case as well, we use Assumption 2.5 on the error distributions under consideration.

Throughout the section, we will use the following set of parameters. Let $m > 4$ be a power of 2, $n = \varphi(m) = m/2$, $d > 1$ an integer, $\Phi_m(X)$ be the m-th cyclotomic polynomial, $\mathbb{K}_m$ be the m-th cyclotomic field, $\mathcal{R}$ its ring of integers, and $\mathcal{M} = \mathcal{R}^d$ the underlying module. Let q be an odd prime such that $q = 1 \mod m$, and let t be an integer such that $2 < t < \sqrt{\log(n)/d}$. The values n and d are the security parameters.

We denote by $\mathcal{M}_q = \mathcal{R}^d / \langle q \rangle$. If $\mathcal{X}$ is a probability distribution over $\mathcal{R}_q$, then $\mathcal{X}^d$ is the distribution over $\mathcal{M}_q$ defined by sampling each coordinate independently from $\mathcal{X}$.

4.1 The MLWE-based public key encryption scheme

Let us generalize the scheme described in Fig. 3.1. The idea is to move from the RLWE setting used in both the key generation and encryption algorithms to the MLWE setting. The scheme is a simplified version of Kyber PKE [9]. Indeed, the structure is the same except for the compression and decompression steps. For this reason, we denote the scheme by KYBER.PKE and define it in Fig. 4.1, where $\mathcal{R}_2$ is the message space, $\mathcal{X}_\sigma$ is a subgaussian probability distribution over $\mathcal{R}_q$, and $\mathcal{X}$ is a t-bounded probability distribution over $\mathcal{R}_q$.

Notice that this time the public key $pk = (\mathbf{A}, B)$ consists of d independent MLWE samples with secret S.

Let $\mu \in \mathcal{R}_2$ be a message. To obtain a decryption failure in KYBER.PKE, the following equation should hold

$$\left\lfloor \langle E, R\rangle - \langle S, E_1\rangle + e_2 + \lfloor \mu \rceil_q \right\rceil_2 \neq \mu.$$

With computations similar to those in Sect. 3.1, we can show that if a decryption failure occurs, then

$$\|\langle E, R\rangle - \langle S, E_1\rangle + e_2\|_\infty \geq (q-2)/4.$$

In particular, this means that there exists an index $k \in \{0, \ldots, n-1\}$ such that

$$\left|\langle \mathbf{v}, \mathbf{w}_{(k)}\rangle + \mathbf{e_2}[k]\right| \geq (q-2)/4, \tag{4.1}$$

where $\mathbf{v}^\top = \left[\mathbf{e}^\top, -\mathbf{s}^\top\right]$ is the concatenation of the coefficient vectors of E and $-S$, and $\mathbf{w}_{(k)}^\top = \left[\mathbf{r}_{(k)}^\top, \mathbf{e_1}_{(k)}^\top\right]$ is the concatenation of the kth rotation coefficient vectors of R and E_1.

Thanks to this condition we can prove that KYBER.PKE is correct with overwhelming probability. This proof is quite similar to the one in Lemma 3.2.

Lemma 4.1 (Correctness) *Let $\mathcal{X}_\sigma$ be a subgaussian distribution, and let $\mathcal{X}$ be a t-bounded probability distribution over $\mathcal{R}_q$. If $q \in \Omega(n\log(n))$ and $\sigma = \Theta\left(\sqrt{\log(n)/d}\right)$, then the PKE scheme described in Fig. 4.1 is correct with overwhelming probability.*

Similarly to Lemma 3.3, we can prove IND-CPA security of KYBER.PKE by applying the hardness of DMLWE twice.

Lemma 4.2 (IND-CPA security) *The PKE scheme described in Fig. 4.1 is* IND-CPA *secure assuming the hardness of* $\mathsf{DMLWE}_{q,\mathcal{X}_\sigma}$ *and* $\mathsf{DMLWE}_{q,\mathcal{X}}$.

4.2 Security reduction with Subgaussian distributions

We follow the same approach as in Sect. 3.2 to prove the reduction from the DMLWE hardness assumption. For this reason, we keep the same assumptions as in Fig. 3.2. We list explicitly in Fig. 4.2 only the modified assumptions needed to account for the module dimension d.

We follow the same idea of Sect. 3.2 to prove the reduction from the DMLWE hardness assumption. For this reason we keep the same assumptions in Fig. 3.2. We explicitly list in Fig. 4.2 only the different assumptions due to the fact that now we have to take into consideration also the degree d of the module.

We are now ready to state the main theorem of this section.

Theorem 4.3 (from DMLWE to FFP-NG). *Let n, d, q, t, $\mathcal{X}$, $\mathcal{X}_\sigma$, $\mathcal{X}_{\sigma'}$, and $\mathcal{X}_{\tilde{\sigma}}$ satisfy the assumptions in Fig. 4.2. For all* FFP-NG *adversaries $\mathcal{A}$ against the* KYBER.PKE *scheme, with errors distribution $\mathcal{X}_\sigma$ and randomness $\mathcal{X}$, denoted by Π, there exists an adversary $\mathcal{B}$ that solves the* DMLWE *problem with error distribution $\mathcal{X}_{\tilde{\sigma}}$ such that*

$$Adv_\Pi^{\mathsf{FFP\text{-}NG}}(\mathcal{A}) \leq Adv_\Pi^{\mathsf{DMLWE}}(\mathcal{B}) + negl(n). \tag{4.2}$$

Assumptions: all parameters satisfy the assumptions in Fig. 3.2 except for the following two changes

- $1 < t < \sqrt{\log(n)/d}$;
- σ, σ', $\hat{\sigma} \in \Theta\left(\sqrt{\log(n)/d}\right)$;

Fig. 4.2. Parameters assumptions for Thm. 4.3.

$\mathcal{B}(\mathbf{A}, B)$:
$S', E' \leftarrow \mathcal{X}_{\sigma'}$
$B' := \mathbf{A}S' + E'$
$pk' := (\mathbf{A}, B + B')$
$(\mu, (R, E_1, e_2)) \leftarrow \mathcal{A}(pk')$
$\mathbf{w}^\top := \left[\mathbf{r}^\top, \mathbf{e_1}^\top\right]$, $\mathbf{v'}^\top := \left[\mathbf{e'}^\top, -\mathbf{s'}^\top\right]$, $\mathbf{w}_{(k)}{}^\top := \left[\mathbf{r}_{(k)}{}^\top, \mathbf{e_1}_{(k)}{}^\top\right]$
$\theta := (q\|\mathbf{w}\|) / (8t\sqrt{n})$
if $\exists k$ s.t. $|\langle \mathbf{v'}, \mathbf{w}_{(k)} \rangle + \mathbf{e_2}[k]| \geq (q - 4\theta - 2)/4$
  return 1 #MLWE
else
  return 0 #Uniform

Fig. 4.3. DMLWE adversary with subgaussian error distribution $\mathcal{X}_\sigma$, and t-bounded randomness distribution $\mathcal{X}$

The DMLWE adversary $\mathcal{B}$ is defined in Fig. 4.3. The proof idea and its structure are identical to those of Thm. 3.4. The main difference is that the kth coefficient of the total error is now expressed as the inner product of two vectors in $\mathbb{R}^{dn}$, rather than in $\mathbb{R}^n$. To prove that $\mathcal{B}$ has a non-negligible advantage at solving DMLWE, we need to show that

$$\left| \Pr_{(\mathbf{A},B)\leftarrow\mathsf{MLWE}}[\mathcal{B}(\mathbf{A}, B) = 1] - \Pr_{(\mathbf{A},B)\leftarrow\$\mathcal{R}_q^{d\times d}\times\mathcal{R}_q^d}[\mathcal{B}(\mathbf{A}, B) = 1] \right| \notin negl(n).$$

Given the tuple (R, E_1, e_2), for $k \in \{0, \ldots, n-1\}$, define the event

$$\mathrm{FAIL}_k(R, E_1, e_2) := \left\{ |\langle \mathbf{v'}, \mathbf{w}_{(k)} \rangle + \mathbf{e_2}[k]| \geq \frac{q - 4\theta - 2}{4} \right\},$$

where $\mathbf{v'}$, $\mathbf{w}_{(k)} \in \mathbb{R}^{dn}$. By definition of $\mathcal{B}$, it suffices to prove that

$$\left| \Pr_{(\mathbf{A},B)\leftarrow\mathsf{MLWE}}[\exists k\text{: } \mathrm{FAIL}_k(R, E_1, e_2)] - \Pr_{(\mathbf{A},B)\leftarrow\$\mathcal{R}_q^{d\times d}\times\mathcal{R}_q^d}[\exists k\text{: } \mathrm{FAIL}_k(R, E_1, e_2)] \right| \tag{4.3}$$

is non-negligible in n. Using the modified parameters described in Fig. 4.2 we can adapt Lemma 3.5 and Lemma 3.6, and bound separately the two terms of Eq. (4.3).

Lemma 4.4 (FAIL with uniformly random key) *Let $\mathbf{v}' \leftarrow \mathcal{X}_{\sigma'}$ be an error vector, and $(R, E_1, e_2) \leftarrow \mathcal{A}(pk')$ be the adversarially chosen encryption randomness. If $(\mathbf{A}, B)$ is sampled uniformly at random, then*

$$\Pr_{(\mathbf{A},B)\leftarrow_{\$}\mathcal{R}_q^{d\times d}\times\mathcal{R}_q^d}[\exists k\colon \mathrm{FAIL}_k(R, E_1, e_2)] \in negl(n). \tag{4.4}$$

Lemma 4.5 (FAIL with MLWE key) *Let $\mathbf{v}' \leftarrow \mathcal{X}_{\sigma'}$ be an error vector, and $(R, E_1, e_2) \leftarrow \mathcal{A}(pk')$ be the adversarially chosen encryption randomness. If $(\mathbf{A}, B)$ is sampled according to the MLWE distribution, then*

$$\Pr_{(\mathbf{A},B)\leftarrow\mathsf{MLWE}}[\exists k\colon \mathrm{FAIL}_k(R, E_1, e_2)] \notin negl(n). \tag{4.5}$$

Combining Lemma 4.4 and Lemma 4.5 proves Thm. 4.3.

A New family of subgaussian

We define the centered binomial distribution ψ_t as follows: given $t \geq 1$ be a fixed value, sample A_t and C_t independently from the binomial distribution $B(t, 1/2)$, and output $A_t - C_t$. By using the properties of the binomial distribution $B(t, 1/2)$ and the fact that A_t and C_t are independent, we get that ψ_t is centered, it has support in $\{-t, \ldots, t\}$, and it is subgaussian with parameter $\sigma_t = \sqrt{2\pi t}$. Using Vandermonde's identity we describe the probability mass function of $X \leftarrow \psi_t$ by

$$\Pr[X = k] = 2^{-2t}\sum_{i=1}^{t}\binom{t}{i}\binom{t}{i-k} = 2^{-2t}\binom{2t}{t+k}, \tag{A.1}$$

for every $k \in \{-t, \ldots, t\}$. Furthermore, the sum of two independent random variables $X \leftarrow \psi_t$ and $Y \leftarrow \psi_s$ is distributed according to ψ_{t+s}. We say that a vector $\mathbf{x}$ (resp. polynomial p) is distributed according to ψ_t if each coordinate $x[i]$ (resp. coefficient p_i) has been sampled independently from ψ_t. We want to show that centered binomials fulfil the defining property of our family of subgaussian. Before doing so, we have to introduce another family of probability distributions. We say that a random variable X follows the hypergeometric distribution $H(N, K, n)$ if its probability mass function is given by

$$\Pr[X = k] = \frac{\binom{K}{k}\binom{N-K}{n-k}}{\binom{N}{n}}, \tag{A.2}$$

Proposition A.1 *Let $X \leftarrow \psi_t$ and $Y \leftarrow \psi_s$ two independent random variables, and let $Z := X + Y$. Consider $z \in \{-(t + s), \ldots, t + s\}$. The variable $X|Z = z$ is identically distributed as $W - t$, where $W \sim H(N, K, n)$ with parameters $N = 2t + 2s$, $K = t + s + z$, and $n = 2t$.*

Proof. Let's analyse the probability mass function of $X|Z=z$. We have

$$\Pr[X=x|Z=z] = \frac{\Pr[X=x]\Pr[Y=z-x]}{\Pr[Z=z]} = \frac{\binom{2t}{t+x}\binom{2s}{s+z-x}}{\binom{2s+2t}{s+t+z}}$$
$$= \frac{\binom{s+t+z}{t+x}\binom{s+t-z}{t-x}}{\binom{2s+2t}{2t}} = \Pr[W=t+x],$$

where W is an hypergeometric random variable with parameters $N=2t+2s$, $K=t+s+z$, and $n=2t$.

The distribution $H=W-t$ has mean $zt/(t+s)$ and variance

$$Var(H) = n\frac{K}{N}\left(1-\frac{K}{N}\right)\frac{N-n}{N-1} = \frac{st}{2s+2t-1}\left(1-\frac{z^2}{(s+t)^2}\right)$$

Using Bernstein inequality and the bounded support, we can show that $H-\mathbb{E}[H]$ is subgaussian with parameter $\sigma_t\sigma_s/\sqrt{\sigma_t^2+\sigma_s}$. This means that centered binomials fulfil the extra condition we used to define our subgaussian family.

We now want to prove that discrete gaussians belongs to our family of subgaussians. Recall that given a lattice Λ and a real σ, the discrete gaussian random variable $X \leftarrow D_{\Lambda,\sigma}$ is defined as

$$\Pr[X=\mathbf{x}] = \frac{\rho_\sigma(\mathbf{x})}{\rho_\sigma(\Lambda)} = \frac{\exp\left(-\pi\frac{\|\mathbf{x}\|^2}{\sigma^2}\right)}{\sum_{\mathbf{y}\in\Lambda}\exp\left(-\pi\frac{\|\mathbf{y}\|^2}{\sigma^2}\right)}$$

Proposition A.2 *Let $X \leftarrow D_{\mathbb{Z}^n,\sigma_1}$ and $Y \leftarrow D_{\mathbb{Z}^n,\sigma_2}$ two independent random variables, let $Z := X+Y$, and let $\sigma^2=\sigma_1^2+\sigma_2^2$ and $\overline{\sigma} := \sigma_1\sigma_2/\sigma$. The variable $X|Z=z$ is distributed according to $D_{\mathbb{Z}^n-\mathbf{z}\sigma_1^2/\sigma^2,\overline{\sigma}}$.*

Proof. Let's write the probability mass function of $X|Z=z$.

$$\Pr[X=\mathbf{x}|Z=\mathbf{z}] = \frac{\Pr[X=\mathbf{x}]\Pr[Y=\mathbf{z}-\mathbf{x}]}{\Pr[Z=\mathbf{z}]} = \frac{\rho_{\sigma_1}(\mathbf{x})\,\rho_{\sigma_2}(\mathbf{z}-\mathbf{x})}{\rho_\sigma(\mathbf{z})\,\rho_{\overline{\sigma}}\left(\mathbb{Z}^n-\frac{\sigma_1^2}{\sigma^2}\mathbf{z}\right)}$$
$$= \frac{1}{\rho_{\overline{\sigma}}\left(\mathbb{Z}^n-\frac{\sigma_1^2}{\sigma^2}\mathbf{z}\right)}\exp\left(-\pi\left(\frac{\|\mathbf{x}\|^2}{\sigma_1^2}+\frac{\|\mathbf{z}-\mathbf{x}\|^2}{\sigma_2^2}-\frac{\|\mathbf{z}\|^2}{\sigma^2}\right)\right)$$

where we used [7, Lemma 4.12] for $\Pr[Z=z]$ and simplified the expression. Now we focus on the exponent

$$\frac{\|\mathbf{x}\|^2}{\sigma_1^2}+\frac{\|\mathbf{z}-\mathbf{x}\|^2}{\sigma_2^2}-\frac{\|\mathbf{z}\|^2}{\sigma^2} = \frac{\|\mathbf{x}-\frac{\sigma_1^2}{\sigma^2}\mathbf{z}\|^2}{\overline{\sigma}^2}$$

which gives us

$$\Pr[X=\mathbf{x}|Z=\mathbf{z}] = \frac{\rho_{\overline{\sigma}}\left(\mathbf{x}-\frac{\sigma_1^2}{\sigma^2}\mathbf{z}\right)}{\rho_{\overline{\sigma}}\left(\mathbb{Z}^n-\frac{\sigma_1^2}{\sigma^2}\mathbf{z}\right)} \sim D_{\mathbb{Z}^n-\frac{\sigma_1^2}{\sigma^2}\mathbf{z},\overline{\sigma}}$$

.

References

[1] Ajtai, M.: Generating hard instances of lattice problems. Electron. Colloquium Comput. Complex. **TR96-007** (1996), https://eccc.weizmann.ac.il/eccc-reports/1996/TR96-007/index.html

[2] Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange - a new hope. IACR Cryptol. ePrint Arch. p. 1092 (2015), http://eprint.iacr.org/2015/1092

[3] Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In: Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings. Lecture Notes in Computer Science, Springer (2009). https://doi.org/10.1007/978-3-642-03356-8_35, https://doi.org/10.1007/978-3-642-03356-8_35

[4] Bindel, N., Hamburg, M., Hülsing, A., Persichetti, E.: Tighter proofs of CCA security in the quantum random oracle model. IACR Cryptol. ePrint Arch. p. 590 (2019), https://eprint.iacr.org/2019/590

[5] Bindel, N., Schanck, J.M.: Decryption failure is more likely after success. IACR Cryptol. ePrint Arch. p. 1392 (2019), https://eprint.iacr.org/2019/1392

[6] Boneh, D., Dagdelen, Ö., Fischlin, M., Lehmann, A., Schaffner, C., Zhandry, M.: Random oracles in a quantum world. In: Lee, D.H., Wang, X. (eds.) Advances in Cryptology – ASIACRYPT 2011. Lecture Notes in Computer Science, vol. 7073, pp. 41–69. Springer, Berlin, Heidelberg, Germany, Seoul, South Korea (Dec 4–8, 2011). https://doi.org/10.1007/978-3-642-25385-0_3

[7] Boneh, D., Freeman, D.M.: Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. IACR Cryptol. ePrint Arch. **2010**, 453 (2010), http://eprint.iacr.org/2010/453

[8] Bos, J.W., Costello, C., Ducas, L., Mironov, I., Naehrig, M., Nikolaenko, V., Raghunathan, A., Stebila, D.: Frodo: Take off the ring! practical, quantum-secure key exchange from LWE. IACR Cryptol. ePrint Arch. p. 659 (2016), http://eprint.iacr.org/2016/659

[9] Bos, J.W., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Stehlé, D.: CRYSTALS - kyber: a cca-secure module-lattice-based KEM. IACR Cryptol. ePrint Arch. p. 634 (2017), http://eprint.iacr.org/2017/634

[10] Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. CoRR **abs/1306.0281** (2013), http://arxiv.org/abs/1306.0281

[11] D'Anvers, J., Batsleer, S.: Multitarget decryption failure attacks and their application to saber and kyber. IACR Cryptol. ePrint Arch. p. 193 (2021), https://eprint.iacr.org/2021/193

[12] D'Anvers, J., Rossi, M., Virdia, F.: (one) failure is not an option: Bootstrapping the search for failures in lattice-based encryption schemes. IACR Cryptol. ePrint Arch. p. 1399 (2019), https://eprint.iacr.org/2019/1399

[13] D'Anvers, J., Rossi, M., Virdia, F.: (one) failure is not an option: Bootstrapping the search for failures in lattice-based encryption schemes. IACR Cryptol. ePrint Arch. p. 1399 (2019), https://eprint.iacr.org/2019/1399

[14] D'Anvers, J., Vercauteren, F., Verbauwhede, I.: On the impact of decryption failures on the security of LWE/LWR based schemes. IACR Cryptol. ePrint Arch. p. 1089 (2018), https://eprint.iacr.org/2018/1089

[15] Dent, A.W.: A designer's guide to kems. IACR Cryptol. ePrint Arch. p. 174 (2002), http://eprint.iacr.org/2002/174

[16] Fujisaki, E., Okamoto, T.: Secure integration of asymmetric and symmetric encryption schemes. J. Cryptol. **26**(1), 80–101 (2013). https://doi.org/10.1007/S00145-011-9114-1, https://doi.org/10.1007/s00145-011-9114-1

[17] Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. IACR Cryptol. ePrint Arch. p. 432 (2007), http://eprint.iacr.org/2007/432

[18] Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the fujisaki-okamoto transformation. IACR Cryptol. ePrint Arch. p. 604 (2017), http://eprint.iacr.org/2017/604

[19] Hofheinz, D., Hövelmanns, K., Kiltz, E.: A modular analysis of the Fujisaki-Okamoto transformation. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017: 15th Theory of Cryptography Conference, Part I. Lecture Notes in Computer Science, vol. 10677, pp. 341–371. Springer, Cham, Switzerland, Baltimore, MD, USA (Nov 12–15, 2017). https://doi.org/10.1007/978-3-319-70500-2_12

[20] Hövelmanns, K., Hülsing, A., Majenz, C.: Failing gracefully: Decryption failures and the Fujisaki-Okamoto transform. In: Agrawal, S., Lin, D. (eds.) Advances in Cryptology – ASIACRYPT 2022, Part IV. Lecture Notes in Computer Science, vol. 13794, pp. 414–443. Springer, Cham, Switzerland, Taipei, Taiwan (Dec 5–9, 2022). https://doi.org/10.1007/978-3-031-22972-5_15

[21] Hövelmanns, K., Kiltz, E., Schäge, S., Unruh, D.: Generic authenticated key exchange in the quantum random oracle model. IACR Cryptol. ePrint Arch. p. 928 (2018), https://eprint.iacr.org/2018/928

[22] Jiang, H., Zhang, Z., Chen, L., Wang, H., Ma, Z.: Ind-cca-secure key encapsulation mechanism in the quantum random oracle model, revisited. In: Shacham, H., Boldyreva, A. (eds.) Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III. Lecture Notes in Computer Science, vol. 10993, pp. 96–125. Springer (2018). https://doi.org/10.1007/978-3-319-96878-0_4, https://doi.org/10.1007/978-3-319-96878-0_4

[23] Jiang, H., Zhang, Z., Ma, Z.: Key encapsulation mechanism with explicit rejection in the quantum random oracle model. IACR Cryptol. ePrint Arch. p. 52 (2019), https://eprint.iacr.org/2019/052
[24] Lang, S.: Algebraic Number Theory, Graduate Texts in Mathematics, vol. 110. Springer-Verlag, New York, 2nd edn. (1994). https://doi.org/10.1007/978-1-4612-0853-2
[25] Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. IACR Cryptol. ePrint Arch. p. 90 (2012), http://eprint.iacr.org/2012/090
[26] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. J. ACM **60**(6), 43:1–43:35 (2013). https://doi.org/10.1145/2535925, https://doi.org/10.1145/2535925
[27] Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-lwe cryptography. IACR Cryptol. ePrint Arch. p. 293 (2013), http://eprint.iacr.org/2013/293
[28] Majenz, C., Sisinni, F.: Provable security against decryption failure attacks from LWE. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology – CRYPTO 2024, Part II. Lecture Notes in Computer Science, vol. 14921, pp. 456–485. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 18–22, 2024). https://doi.org/10.1007/978-3-031-68379-4_14
[29] National Institute of Standards and Technology: Module-lattice-based key-encapsulation mechanism standard (fips 203). Federal Information Processing Standard 203, U.S. Department of Commerce, National Institute of Standards and Technology (aug 2024), https://doi.org/10.6028/NIST.FIPS.203
[30] Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem. Electron. Colloquium Comput. Complex. **TR08-100** (2008), https://eccc.weizmann.ac.il/eccc-reports/2008/TR08-100/index.html
[31] Peikert, C.: A decade of lattice cryptography. IACR Cryptol. ePrint Arch. p. 939 (2015), http://eprint.iacr.org/2015/939
[32] Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. IACR Cryptol. ePrint Arch. p. 348 (2007), http://eprint.iacr.org/2007/348
[33] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM **56**(6), 34:1–34:40 (2009). https://doi.org/10.1145/1568318.1568324, https://doi.org/10.1145/1568318.1568324
[34] Saito, T., Xagawa, K., Yamakawa, T.: Tightly-secure key-encapsulation mechanism in the quantum random oracle model. IACR Cryptol. ePrint Arch. p. 1005 (2017), http://eprint.iacr.org/2017/1005
[35] Stehlé, D., Steinfeld, R.: Making ntruencrypt and ntrusign as secure as standard worst-case problems over ideal lattices. IACR Cryptol. ePrint Arch. p. 4 (2013), http://eprint.iacr.org/2013/004
[36] Vershynin, R.: High-Dimensional Probability: An Introduction with Applications in Data Science. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambrige, 2nd edn. (2026)