Provable security against decryption failure attacks from LWE

Christian Majenz[0000−0002−1877−8385], Fabrizio Sisinni[0009−0007−9641−4329]

Technical University of Denmark

Abstract. In a recent work, Hövelmanns, Hülsing and Majenz introduced a new security proof for the Fujisaki-Okamoto transform in the quantum-accessible random oracle model (QROM) used in post-quantum key encapsulation mechanisms. While having a smaller security loss due to decryption failures present in many constructions, it requires two new security properties of the underlying public-key encryption scheme (PKE).
In this work, we show that one of the properties, *Find Failing Plaintexts - Non Generic* (FFP-NG) security, is achievable using a relatively efficient LWE-based PKE that does not have perfect correctness. In particular, we show that LWE reduces to breaking FFP-NG security of the PVW scheme, when all LWE errors are discrete Gaussian distributed. The reduction has an arbitrarily small constant multiplicative loss in LWE error size. For the proof, we make use of techniques by Genise, Micciancio, Peikert and Walter to analyze marginal and conditional distributions of sums of discrete Gaussians.

1 Introduction

At least since the launch of the NIST standardization effort for post-quantum (PQ) key encapsulation mechanisms (KEMs) and digital signature schemes (DSS') [21], there has been an increasing focus on post-quantum cryptography. The Fujisaki-Okamoto (FO) transform has become the go-to technique to upgrade a basic post-quantum public-key encryption scheme (PKE) to a chosen-ciphertext-secure KEM. In fact, all candidate KEMs that have survived the first round of scrutiny use a version of this transformation, including ML-KEM [25]. It was therefore imperative to provide a proof of security for the used variants of the FO transformation that is as tight as possible. The line of work striving to provide such a proof in the quantum-accessible random oracle model (QROM) started with [26,13] and has seen many improvements since [24,17,7,15,18,11,14].

Many of the most efficient post-quantum PKEs have a small probability that decryption of an honestly generated ciphertext fails, i.e., they only enjoy *approximate* and not *perfect* correctness. Examples include the PKE underlying ML-KEM/Kyber [25,9], and the PKEs underlying BIKE [1] and HQC [19] which are still under consideration in round four of the NIST process. It is therefore important that security proofs for the FO transformation allow for approximately correct PKEs. One of the most recent works on the FO transformation [14]

provides a blueprint for improving the way decryption failures are handled in the security reduction for the FO transformation.[1] To that end, [14] introduced a family of security notions, the Find Failing Plaintext (FFP) notions. Security for two members of this family, Find Failing Plaintexts - Non Generic (FFP-NG)[2] and Find Failing Plaintexts - No Key (FFP-NK), are prerequisites for the results of [14] to be applicable, the alternative being to revert to a manifestly non-tight analysis of decryption failures [16]. While FFP-NK is statistical in nature and can be taken care of by characterizing mean and variance of a decryption-error-related probability distribution [14], FFP-NG will rely on a computational assumption for PKEs with approximate correctness. As a consequence, FFP-NG security should be analyzed for any approximately correct post-quantum PKE to ensure applicability of the tighter analysis of decryption failures from [14], through cryptanalysis or a security reduction to a well-studied quantum-hard problem. It is important to note that FFP-NG security is independent of IND-CPA security (see Section 3).

In this work, we give the first result in this direction: we prove the FFP-NG security of the learning-with-errors-based (LWE-based) PKE introduced by Peikert, Vaikuntanathan and Waters (the PVW PKE) assuming the hardness of the LWE problem, when all LWE errors are drawn from a discrete Gaussian distribution. The PVW PKE is a generalization of Regev's original LWE-based PKE [23] with improved efficiency. We use discrete Gaussians instead of rounded Gaussians since they have better algebraic properties. In particular, under some assumptions on the smoothing parameters of the lattices the random variables are supported on, convolutions and marginal probabilities can be easily analyzed as they are approximately distributed according to discrete Gaussian distributions as well. We use the modular framework introduced by Genise, Micciancio, Peikert and Walter in [12] to make efficient use of this fact.

Our Result. We go on to give an informal description of our main result. We prove the following security reduction.

Theorem 1 (Decision LWE $\Longrightarrow$ FFP-NG, Informal). *Let $\mathcal{A}$ be an FFP-NG adversary against the PVW encryption scheme with discrete Gaussian LWE errors with standard deviation σ, denoted by Π. If the FFP-NG advantage of $\mathcal{A}$ is non-negligible in the number of LWE secrets n in the secret key of Π, there exists an adversary $\mathcal{B}$ that solves the Decision LWE problem with discrete Gaussian error distribution with standard deviation $\hat{\sigma} = \sigma/\varphi$ for any constant $\varphi > 1$.*

Technical Overview. We provide an informal technical overview. To that end, we need to describe the FFP-NG security game. In the FFP-NG game for a PKE

[1] The reduction presented in [14] only applies if the underlying PKE is genuinely randomized, i.e. if the ciphertexts have min-entropy ("gamma-spreadness").

[2] In [14] the security game is called Find Failing Plaintext that are Non Generic (FFP-NG). Since the "Non Generic" refers to the fact that these plaintexts fail only for specific keys, we change the name of the game keeping the acronym equal to stress belonging to the FFP family.

Π two key pairs (sk_0, pk_0) and (sk_1, pk_1) for Π are generated. The adversary receives a public key pk_0 as input. The adversary can then submit a message together with a random tape for the (randomized) encryption algorithm. Now a bit $b \leftarrow \{0,1\}$ is drawn at random, and the message is encrypted by using public key pk_b and the random tape provided by the adversary. Subsequently, the resulting ciphertext is decrypted with key sk_b. If decryption fails to return the original message, the adversary receives answer 1, otherwise answer 0. Finally the adversary has to submit a guess b' of b and wins if it is correct.

For the PVW PKE, the decryption error probability for any fixed pair of message and randomness and a freshly generated key pair like (pk_1, sk_1) is negligible. Therefore the only way to win the FFP-NG game is to force a decryption failure for the key pair (pk_0, sk_0).

For the PVW PKE, the encryption randomness determines the coefficients $\mathbf{d}$ of the linear combination of the LWE samples from the public key that is used to hide the message. To force PVW decryption to fail, intuitively speaking, an adversary has to choose these coefficients such that the LWE errors build up and do not cancel. Geometrically, the encryption randomness vector needs to be aligned with the error vector, i.e. they need to have a large inner product.

The basic idea of our reduction is as follows. The reduction, trying to determine whether the input $\mathbf{s}$ is an LWE sample or uniformly random, adds a small additional error $\mathbf{e}'$ to the input. In case the input is an LWE sample, the result $\mathbf{s}'$ is an LWE sample with a slightly larger LWE error. In case the input it uniformly random, the result is uniformly random and independent of $\mathbf{e}'$. Now the reduction uses $\mathbf{s}'$ to construct a public key for PVW and simulates the FFP-NG game for the assumed successful adversary $\mathcal{A}$. The crucial observation is now as follows.

- In case $\mathbf{s}$ is an LWE sample, $\mathbf{s}'$ is an LWE sample with an error vector that is *correlated* with $\mathbf{e}'$. As we have observed above, to win FFP-NG, $\mathcal{A}$ needs to align its encryption randomness vector $\mathbf{d}$ with the error vector. This implies that $\mathbf{d}$ will be somewhat aligned with $\mathbf{e}'$.
- In case $\mathbf{s}$ is uniform, $\mathbf{s}'$ and thus $\mathbf{d}$ are independent of $\mathbf{e}'$, and thus not very aligned.

The reduction can therefore detect whether $\mathbf{d}$ and $\mathbf{e}'$ are significantly more aligned than expected for independent random ones, and use the result to decide the LWE problem. The main technical work lies in finding an appropriate threshold for this decision criterion, and analyzing the joint distribution of the different LWE errors to prove that the resulting decision LWE algorithm has non-negligible advantage.

Additional Related Work. The idea of adding an additional error to an LWE sample is used in the context of *noise flooding* [6]. In noise flooding, a much larger error is added to an existing LWE sample to hide the original LWE error. This technique is used whenever fine-grained decryption privileges for LWE-based PKE are needed, e.g. for threshold cryptography [6,4,8,10] or electronic voting

[3,2]. Our technique adds an error to an LWE sample with the very different goal of recognizing that error at a later stage, rather than hiding the original error.

2 Preliminaries

In this section we include some notation and introduce some results that will be helpful throughout the rest of the paper. In the rest of the paper we use lowercase bold letters to denote vectors, capital bold letters to denote matrices, and calligraphic capital letters to denote probability distributions.

2.1 Lattices

An n-dimensional lattice Λ is a discrete subgroup of $\mathbb{R}^n$. Given a full rank matrix $\mathbf{B} \in \mathbb{R}^{n \times k}$, we denote the lattice generate by $\mathbf{B}$ as $\mathcal{L}(\mathbf{B}) = \{\mathbf{Bv} \mid \mathbf{v} \in \mathbb{Z}^k\}$. We define the rank of the lattice $\mathcal{L}(\mathbf{B})$ as the rank of the matrix $\mathbf{B}$. We say that the lattice is full-rank if $n = k$, or, equivalently if $\mathbf{B}$ is invertible. We call $\mathbf{B}$ a basis of the lattice. We can easily see that the basis $\mathbf{B}$ is not unique. Indeed, for every unitary matrix $\mathbf{U} \in \mathbb{Z}^{k \times k}$, $\mathbf{B'} = \mathbf{BU}$ is another basis of the same lattice. For every lattice Λ we can define another lattice, called the dual lattice, as $\Lambda^* = \{\mathbf{v} \in span(\Lambda) \mid \langle \mathbf{v}, \Lambda \rangle \in \mathbb{Z}\}$. Furthermore, if $\mathbf{B}$ is a basis of Λ, we have that $\mathbf{B}^{-T}$ is a basis of Λ^*. Let $\Lambda \subset \mathbb{R}^n$ be a lattice, we define a coset of Λ as a set of the form $\Lambda + \mathbf{v}$ for any $\mathbf{v} \in \mathbb{R}^n$. We define a Λ-subspace as the linear span of some set of lattice points, equivalently we can say that L is Λ-subspace if it is a subspace and $L = span(\Lambda \cap L)$.

2.2 Probability notations

We introduce notations similar to those used in [12]. For any probability distribution $\mathcal{X}$ over a set X, a predicate P on X and a function $f\colon X \to Y$ we denote with $[\![f(x) \mid x \leftarrow \mathcal{X}, P(x)]\!]$ for the probability distribution over Y obtained by sampling x according to $\mathcal{X}$, conditioning on x satisfying P, and outputting $f(x) \in Y$. We write $[P(x) \mid x \leftarrow \mathcal{X}]$ to denote the event that $P(x)$ is satisfied when x is selected according to $\mathcal{X}$, and use $\Pr[z \leftarrow \mathcal{X}]$ as an abbreviation for $\mathcal{X}(z) = \Pr[x = z \mid x \leftarrow \mathcal{X}]$. We write $f(\mathcal{X}) = [\![f(x) \mid x \leftarrow \mathcal{X}]\!]$ for the result of applying a function to a probability distribution. We let U_X denote the uniform distribution over a finite set X and we write only U if the set is clear from the context. Let $\mathcal{X}$ and $\mathcal{Y}$ be two probability distributions. We define the statistical distance between $\mathcal{X}$ and $\mathcal{Y}$ such as

$$\Delta(\mathcal{X}, \mathcal{Y}) = \sup_A \left| \Pr[x \in A \mid x \leftarrow \mathcal{X}] - \Pr[y \in A \mid y \leftarrow \mathcal{Y}] \right|,$$

where the supremum is computed over all the measurable sets A. We also define the max-log distance such as

$$\Delta_{ML}(\mathcal{X}, \mathcal{Y}) = \sup_A \left| \log(\Pr[x \in A \mid x \leftarrow \mathcal{X}]) - \log(\Pr[y \in A \mid x \leftarrow \mathcal{Y}]) \right|,$$

where, again, the supremum is taken over all the measurable sets A.
Let $\mathcal{X}$, $\mathcal{Y}$ be two probability distributions defined over the same set and let $\varepsilon \in (0,1)$. We write $\mathcal{X} \overset{\varepsilon}{\approx} \mathcal{Y}$ if, for every z, $\mathcal{X}(z) \in [1-\varepsilon, 1+\varepsilon] \cdot \mathcal{Y}(z)$ and $\mathcal{Y}(z) \in [1-\varepsilon, 1+\varepsilon] \cdot \mathcal{X}(z)$. In this case, if x is distributed according to $\mathcal{X}$, we will say that x is ε-distributed according to $\mathcal{Y}$.
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 $f \in negl(n)$.

2.3 Discrete Gaussian

For any positive definite matrix $\Sigma \in \mathcal{M}(n, \mathbb{R})$ and any vector $\mathbf{x} \in \mathbb{R}^n$, we denote by

$$\rho_{\sqrt{\Sigma}}(\mathbf{x}) = \mathrm{e}^{-\pi \mathbf{x}^T \sqrt{\Sigma}^T \sqrt{\Sigma} \mathbf{x}},$$

the probability density function of an n-dimensional Gaussian distribution with mean 0 and covariance matrix Σ. If the covariance matrix is the identity matrix $\mathbf{I}$, we simply write

$$\rho(\mathbf{x}) = \mathrm{e}^{-\pi \|\mathbf{x}\|^2}.$$

We denote with $\mathcal{D}_{A,\sqrt{\Sigma}}$ the Gaussian distribution over the set A with covariance matrix Σ. In particular, if $\Sigma = \sigma^2 \mathbf{I}$ is a scalar matrix, we just write $\mathcal{D}_{A,\sigma}$. For any lattice Λ and any lattice coset $A \subset \Lambda$, we write

$$\rho_{\sqrt{\Sigma}}(A) = \sum_{\mathbf{x}\in A} \rho_{\sqrt{\Sigma}}(\mathbf{x}).$$

We define the discrete Gaussian distribution over the coset A with covariance matrix Σ as

$$\mathcal{D}_{A,\sqrt{\Sigma}} = \frac{\rho_{\sqrt{\Sigma}}(\mathbf{x})}{\rho_{\sqrt{\Sigma}}(A)}.$$

For every value $\mathbf{z} \in \mathbb{R}^n$ and a lattice Λ, we can also define the $\mathbf{z}$-centered discrete Gaussian distribution as $\mathbf{z} + \mathcal{D}_{\Lambda-\mathbf{z},\sqrt{\Sigma}}$. We have introduced a probability distribution over lattices and the following result is a tail bound for this distribution.

Lemma 1 (Tail Bound, [5] Lemma 1.5). *For any $c > 1/\sqrt{2\pi}$ and an n-dimensional lattice Λ, we have the following bound*

$$\rho\left(\Lambda \setminus c\sqrt{n} B^n(0,1)\right) < C^n \rho(\Lambda), \tag{1}$$

where $C = c \cdot \sqrt{2\pi \mathrm{e}} \cdot \mathrm{e}^{-\pi c^2} < 1$ and $B^n(0,1)$ is the n-dimensional ball centered in 0 with radius 1.

By using the definition of discrete Gaussian $\mathcal{D}_\Lambda$, we can rewrite the bound as

$$\Pr\left[|Z| > c\sqrt{n}\right] < C^n,$$

where Λ is an n-dimensional lattice, $Z \leftarrow \mathcal{D}_\Lambda$, and $C = c \cdot \sqrt{2\pi \mathrm{e}} \cdot \mathrm{e}^{-\pi c^2}$.
Now we define an important tool to study discrete Gaussian distributions, the smoothing parameter.

Definition 1. *For a lattice Λ and any $\varepsilon > 0$, we define the **smoothing parameter** $\eta_\varepsilon(\Lambda)$ to be the smallest value σ such that $\rho_{1/\sigma}(\Lambda^* \setminus \{0\}) \leq \varepsilon$.*

In other words, $\eta_\varepsilon(\Lambda)$ is the smallest σ such that a Gaussian measure scaled by $1/\sigma$ on the dual lattice gives all but $\varepsilon/(1+\varepsilon)$ of its weight to the origin. Observe that the smoothing parameter just defined is a scalar. We can extend this definition to matrices. For any positive definite matrix Σ, we say that $\eta_\varepsilon(\Lambda) \leq \sqrt{\Sigma}$, if $\eta_\varepsilon\left(\sqrt{\Sigma}^{-1} \cdot \Lambda\right) \leq 1$. Equivalently, we say that $\eta_\varepsilon(\Lambda) \leq \sqrt{\Sigma}$ if $\rho\left(\sqrt{\Sigma}^T \cdot \Lambda^* \setminus \{0\}\right) = \rho_{\sqrt{\Sigma^{-1}}}(\Lambda^* \setminus \{0\}) \leq \varepsilon$. We describe a useful bound for the smoothing parameter in the following lemma.

Lemma 2 ([20] Lemma 3.3). *For any n-dimensional lattice Λ and any $\varepsilon > 0$, we have*

$$\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln\left(2n\left(1+1/\varepsilon\right)\right)}{\pi}} \cdot \lambda_n(\Lambda),$$

where $\lambda_n(\Lambda)$ is the minimum length of a set of n linearly independent vectors from Λ, that is the length of the longest vector in the set.
In particular, for $\Lambda = \mathbb{Z}^n$, we have

$$\eta_\varepsilon(\mathbb{Z}^n) \leq \sqrt{\frac{\ln\left(2n\left(1+1/\varepsilon\right)\right)}{\pi}}. \tag{2}$$

Following again [12], we describe the convolution of two random variables $\mathcal{X}$ and $\mathcal{Y}$, defined over X and Y respectively, by using the statistical experiment where

$$x \leftarrow \mathcal{X}\ , y \leftarrow x + \mathcal{Y}_{Y-x}.$$

We denote the resulting probability distribution simply with $x + \mathcal{Y}_{Y-x}$, every time the distribution of x is clear from the context. The next two theorems provide properties of the sum of Discrete Gaussians and the conditioned probability distributions. In the Chapter 3 we are going to use them often.

Theorem 2 (Sum of Gaussians and marginal distribution, [12] Theorem 4.5). *Let $\varepsilon \in (0,1)$, define $\bar{\varepsilon} = 2\varepsilon/(1-\varepsilon)$ and $\varepsilon' = 4\varepsilon/(1-\varepsilon)^2$, let A_1, A_2 be cosets of full-rank lattices Λ_1, Λ_2, let Σ_1, Σ_2 be positive definite matrices where $\eta_\varepsilon(\Lambda_2) \leq \sqrt{\Sigma_2}$, and let*

$$\chi = \left[\!\!\left[(x_1, x_2) \,\middle|\, x_1 \leftarrow \mathcal{D}_{A_1,\sqrt{\Sigma_1}}, x_2 \leftarrow x_1 + \mathcal{D}_{A_2-x_1,\sqrt{\Sigma_2}}\right]\!\!\right].$$

If $\eta_\varepsilon(\Lambda_1) \leq \sqrt{\Sigma_3}$, *where* $\Sigma_3^{-1} = \Sigma_1^{-1} + \Sigma_2^{-1}$ *is a positive definite matrix, then the marginal distribution* χ_2 *of* x_2 *in* χ *satisfies*

$$\chi_2 \overset{\varepsilon'}{\approx} \mathcal{D}_{A_2, \sqrt{\Sigma_1 + \Sigma_2}}.$$

In any case, the distribution $\chi_1^{x_2}$ *of* x_1 *conditioned on any* $x_2 \in A_2$ *satisfies*

$$\chi_1^{x_2} \overset{\bar{\varepsilon}}{\approx} x_2' + \mathcal{D}_{A_2 - x_2', \sqrt{\Sigma_3}},$$

where $x_2' = \Sigma_3 \Sigma_2^{-1} x_2$.

Theorem 3 (Linear combination of Gaussians, [12] Theorem 4.6). *Let* $\varepsilon \in (0,1)$, *let* $\mathbf{z} \in \mathbb{Z}^m \setminus \{0\}$, *and for* $i = 1, \ldots, m$ *let* Λ_i *be* n*-dimensional lattice such that* $\Lambda_\cap = \bigcap_{i=1}^m \Lambda_i$ *is full rank. Let further* $A_i = \Lambda_i + a_i \subset \mathbb{R}^n$ *be a lattice coset and* $S_i \in \mathbb{R}^{n \times n}$ *such that* $A_i \subset \mathrm{Span}_i(S_i)$. *If* $\eta_\varepsilon\left(\ker\left(\mathbf{z}^T \otimes I_n\right) \cap \Lambda\right) \leq S$, *where* $\Lambda = \Lambda_1 \times \cdots \times \Lambda_m$ *and* $S = \mathrm{diag}(S_1, \ldots, S_m)$ *then*

$$\Delta_{ML}\left(\sum_{i=1}^m z_i \mathcal{D}_{A_i, S_i}, \mathcal{D}_{A', S'}\right) \leq \log \frac{1+\varepsilon}{1-\varepsilon},$$

where $A' = \sum_{i=1}^m z_i A_i$ *and* $S' = \sqrt{\sum_{i=1}^m z_i^2 S_i S_i^T}$.
In particular, let each $S_i = s_i I_n$ *for some* $s_i > 0$, *if*

$$\left(\left(\frac{z_j}{s_j}\right)^2 + \max_{i \neq j}\left(\frac{z_i}{s_i}\right)^2\right)^{-1/2} \geq \eta_\varepsilon(\Lambda_\cap)$$

where j *minimizes* $|z_j/s_j| \neq 0$. *Then*

$$\Delta_{ML}\left(\sum_{i=1}^m z_i \mathcal{D}_{A_i, s_i}, \mathcal{D}_{A', s'}\right) \leq \log \frac{1+\varepsilon'}{1-\varepsilon'},$$

where $s' = \sqrt{\sum_{i=1}^m (z_i s_i)^2}$ *and* $1 + \varepsilon' = (1+\varepsilon)^m$.

Observe that, in the theorem above we make use of the isomorphism between $(\mathbb{R}^n)^m$ and $\mathbb{R}^n \otimes \mathbb{R}^m$, with $\mathbb{Z}^m \subset \mathbb{R}^m$ naturally embedded.

2.4 Learning With Errors (LWE)

The Learning With Errors problem was first introduced by Regev in [23] and it is one of the most promising problems used to build post-quantum cryptographic protocols. Let p and n be integers, $\mathcal{X}$ a probability distribution on $\mathbb{Z}_p$ and $\mathbf{s} \in \mathbb{Z}_p^n$ be a vector. We define a probability distribution $A_{\mathbf{s},\mathcal{X}}$ on $\mathbb{Z}_p^n \times \mathbb{Z}_p$ obtained by choosing uniformly at random a vector $\mathbf{a} \in \mathbb{Z}_p^n$, sampling an error $e \in \mathbb{Z}_p$ according to $\mathcal{X}$ and outputting the pair $(\mathbf{a}, b)$, where $b = \langle \mathbf{a}, \mathbf{s} \rangle + e$.

Definition 2. *Let p be an integer and $\mathcal{X}$ a probability distribution on $\mathbb{Z}_p$. Given a pair $(\mathbf{a}, b) \in \mathbb{Z}_p^n \times \mathbb{Z}_p$, we define the* ***Decision LWE*** *problem with modulus p and error distribution $\mathcal{X}$ as the problem of determine if the given pair has been sampled according to $A_{\mathbf{s},\mathcal{X}}$ or uniformly at random from $\mathbb{Z}_p^n \times \mathbb{Z}_p$.*

We define a public key encryption scheme based on the LWE problem. The scheme was first introduced by Regev in [23] and then improved by Peikert, Vaikuntanathan and Waters in [22].
We present the latter scheme with discrete Gaussian as error distribution, called $PVW_{\mathcal{D}}$, in Algorithm 1. Our notation is based on the presentation of the PVW scheme in [20].

Algorithm 1 PVW PKE [22,20] with Discrete Gaussian error distribution

- **Parameters:** Integers n, m, ℓ, t, r, p and a real number $\alpha > 0$.
- **Private Key:** Choose $\mathbf{S} \in \mathbb{Z}_p^{n \times \ell}$ uniformly at random. The private key is $\mathbf{S}$.
- **Public Key:** Choose $\mathbf{A} \in \mathbb{Z}_p^{n \times m}$ uniformly at random and $\mathbf{E} \in \mathbb{Z}_p^{\ell \times m}$ by choosing each entry according to $\mathcal{D}_{\mathbb{Z}_p,\sigma}$. Set $\mathbf{B} = \mathbf{S}^T\mathbf{A} + \mathbf{E}$. The public key is $(\mathbf{A}, \mathbf{B}) \in \mathbb{Z}_p^{n \times m} \times \mathbb{Z}_p^{\ell \times m}$.
- **Encryption:** Given an element of the message space $\mathbf{v} \in \mathbb{Z}_t^\ell$ and a public key $(\mathbf{A}, \mathbf{B})$, choose a vector $\mathbf{d} \in \{-r, \ldots, r\}^m$ uniformly at random. Set $\mathbf{u} = \mathbf{A}\mathbf{d}$ and $\mathbf{c} = \mathbf{B}\mathbf{d} + \lfloor (p/t)\,\mathbf{v} \rceil$. The ciphertext is $(\mathbf{u}, \mathbf{c}) \in \mathbb{Z}_p^n \times \mathbb{Z}_p^\ell$.
- **Decryption:** Given a ciphertext $(\mathbf{u}, \mathbf{c}) \in \mathbb{Z}_p^n \times \mathbb{Z}_p^\ell$ and a private key $\mathbf{S} \in \mathbb{Z}_p^{n \times \ell}$, output $\lfloor (t/p)\left(\mathbf{c} - \mathbf{S}^T\mathbf{u}\right) \rceil$.

2.5 Find Failing Plaintext - Non Generic (FFP-NG)

In [14] the authors introduce a novel framework to analyze the impact of decryption failures on the security of Public Key Encryption schemes. They introduce the family of security games **Find Failing Plaintext (FFP)**. We are mainly interested in one member of the family, Find Failing Plaintext - Non Generic (FFP-NG). In this game an adversary $\mathcal{A}$ is provided with a public key and should find a message-randomness pair that triggers a decryption failure more likely with the respect to the key that $\mathcal{A}$ sees than with respect to an independent key pair. In [13], Hofheinz et al. introduce a correctness game, called COR, that aims to describe how to handle with decryption failure. This game looks unnatural in most contexts, due to the fact that the game provides the adversary a secret key. In contrast to the COR game, the FFP-NG game provides the adversary only a public key and one query to the Failure Checking Oracle (FCO). This game allows to analyze the hardness of finding meaningful decryption failures independently from the hardness of searching a random oracle for them. We define the FFP-NG game as shown in the Algorithm 2. In this game, the adversary $\mathcal{A}$ gets a public key pk_0 and it is allowed to query a single message-randomness pair to the Failure Checking Oracle, that is given

Algorithm 2 FFP-NG game against a PKE $\Pi = (\mathrm{KeyGen}, \mathrm{Enc}, \mathrm{Dec})$

$(sk_0, pk_0) \leftarrow \mathrm{KeyGen}$
$(sk_1, pk_1) \leftarrow \mathrm{KeyGen}$
$b \leftarrow \{0,1\}$
$b' \leftarrow \mathcal{A}^{\mathrm{FCO}_b}(pk_0)$
return $[\![b = b']\!]$
$\underline{\mathbf{FCO}_b(v;d)}$ ▷ One query
$c \leftarrow \mathrm{Enc}(pk_b, v; d)$
$v' := Dec(sk_b, c)$
return $[\![v \neq v']\!]$

access to the key pair (sk_b, pk_b). We want to highlight that the two key pairs are independent from each other. By using this game we can define the advantage of the adversary $\mathcal{A}$ against a PKE scheme $\Pi = (\mathrm{KeyGen}, \mathrm{Enc}, \mathrm{Dec})$ as

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

3 Initial Characterization of FFP-NG

In this section we describe some more or less straightforward properties of FFP-NG security. The results in this section are independent of the remainder of the paper and are meant to record a more complete characterization picture for FFP-NG, and motivate the result in Section 4.

We begin by observing that the notion is trivially achievable.

Proposition 1. *Any PKE that has perfect correctness is FFP-NG secure.*

While it is calming to note that the security definition is not vacuous, the above proposition is also useless, as the entire motivation for the introduction of FFP-NG security is dealing with PKE *without* perfect correctness.

Another important observation is that FFP-NG security is indeed independent from the standard security notions for PKE, i.e. IND-CPA and OW-CPA.[3] This is quite intuitive as well, as the FFP-NG crucially allows the adversary to control the encryption randomness, which the adversary has no control over in the IND-CPA or OW-CPA security games.

Proposition 2. *1. There exists a PKE Π that is OW-CPA-insecure but FFP-NG-secure.*
2. Assuming the existence of IND-CPA-secure PKE, there exists a PKE Π' that is IND-CPA-secure but FFP-NG-insecure.

Proof. The first part follows from the existence of OW-CPA-insecure perfectly correct PKE and Proposition 1. To see the second part, let $\tilde{\Pi} = (\mathsf{Gen}, \widetilde{\mathsf{Enc}}, \mathsf{Dec})$

[3] In this short section we use standard definitions of IND-CPA and OW-CPA. As they only appear in this section which stands independent from the main results of the paper, we refrain from stating these definitions.

be an IND-CPA-secure, perfectly correct PKE. Without loss of generality, assume that $\widetilde{\mathsf{Enc}}$ uses a random tape r that is longer then $\tilde{\Pi}$'s public keys. Now define

$$\mathsf{Enc}_{pk}(m;r) = \begin{cases} \widetilde{\mathsf{Enc}}_{pk}(\bar{m};r) & \text{if } pk \text{ is a prefix of } r \\ \widetilde{\mathsf{Enc}}_{pk}(m;r) & \text{else,} \end{cases}$$

and set $\Pi' = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$. Here, $\bar{m} = m \oplus 1 \dots 1$ is m with all bits flipped. The IND-CPA security of $\tilde{\Pi}$ implies in particular that $\tilde{\Pi}$ is secure against key recovery. Therefore the length of pk needs to be superlogarithmic in the security parameter. It follows that Π' is still IND-CPA secure despite the "puncture", as the chance of the first case being active in any given IND-CPA game is negligible. Π' is, on the other hand, not FFP-NG secure due to the following simple adversary: On input pk, submit the all-zero message and pk padded with zeroes as randomness to the failure checking oracle FCO. Then output the bit received from the FCO.

□

4 Reduction from LWE to FFP-NG

In this section we describe a security reduction from the LWE problem to winning the FFP-NG game against the $PVW_{\mathcal{D}}$ scheme described in Algorithm 1. We recall that the error for all LWE samples is sampled from discrete Gaussians.
We begin by building some intuition. We show that an adversary needs to make use of the input public key.

Proposition 3. *Let Π be the $PVW_{\mathcal{D}}$ encryption scheme from Algorithm 1 and let $\mathcal{A}$ be an FFP-NG adversary against Π. If $\mathcal{A}$ chooses the message-randomness pair $(v, \mathbf{d})$ independent of the given key pk_0 then*

$$Adv^{FFP\text{-}NG}_{\mathcal{A},\Pi} = 0.$$

Proof. Let's recall that by definition the advantage of $\mathcal{A}$ is

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

Thus, we can equivalently prove that $\Pr\left[\text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1\right] = 1/2$. If $\mathcal{A}$ chooses $(v, \mathbf{d})$ independent of the given public key pk_0, the probability of decryption failure is the same no matter $\mathcal{A}$ uses FCO_0 or FCO_1, that is

$$\Pr\left[\mathrm{FCO}_b\left(v, \mathbf{d}\right) = 1 | b = 0\right] = \Pr\left[\mathrm{FCO}_b\left(v, \mathbf{d}\right) = 1\right] = \Pr\left[\mathrm{FCO}_b\left(v, \mathbf{d}\right) = 1 | b = 1\right],$$

$$\Pr\left[\mathrm{FCO}_b\left(v, \mathbf{d}\right) = 0 | b = 0\right] = \Pr\left[\mathrm{FCO}_b\left(v, \mathbf{d}\right) = 0\right] = \Pr\left[\mathrm{FCO}_b\left(v, \mathbf{d}\right) = 0 | b = 1\right].$$

Thanks to these equations we can say that each input $\mathcal{A}$ gets is independent of b. Thus, also the output of $\mathcal{A}$ is independent of b, that is $\Pr\left[\text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1\right] = 1/2$.

□

In words, the adversary has to use the public key it receives as input in order to win the game. As part of the proof of the main theorem, we will show that the only strategy the adversary can use is to find a message-randomness pair that triggers a decryption failure with higher probability with respect to the input key than with respect to an independent key pair.

At a high level, the security reduction works as follows. Let $\mathcal{A}$ be an attacker that wins the FFP-NG game for $PVW_{\mathcal{D}}$ for some error size σ with noticeable probability. We construct an algorithm that solves the decision LWE problem for a slightly smaller error size $\hat{\sigma}$. The algorithm samples an additional error with size $\sigma' = \sqrt{\sigma^2 - \hat{\sigma}^2}$ and adds it to its input s which is either an LWE sample ("real") or uniform ("random"). Let the result be s'. In the real case, s' is almost exactly an LWE sample with error size σ. In the random case, on the other hand, adding an independently sampled error to a uniformly random value just yields a uniformly random value. The purpose of adding additional noise is, in some sense, opposite to the so-called noise flooding technique [6]: while noise flooding has the goal of hiding the original error, the goal of adding an additional error here is to obtain partial knowledge about the error. The algorithm now simulates the adversary $\mathcal{A}$, using s' as public key. It will then *correlate* the encryption randomness the adversary submits to the failure checking oracle with the added noise. In the random case, the noise and encryption randomness are manifestly independent and thus uncorrelated. In the real case, on the other hand, the adversary will likely cause a decryption failure by the assumption that it wins the FFP-NG game. This, however, is only possible if the encryption randomness is aligned with the LWE error, which means it is correlated with the added noise.

The main result is given in Theorem 4. Unless stated differently, we assume the conditions on parameters $p, m, \ell, r, t, \alpha, \sigma$ of the LWE scheme and the constant φ given in Table 1. To work with column vectors instead of row vectors,

Conditions on parameters
$r \geq 1$
$t \geq 2$
$\ell = 1$
$m = n^c$, with $1 < c < 2$
$p \geq 2mtr \ln(n)$, with p a prime number
$1 < \varphi < \min\{\sqrt{m}/10, m^{1/4}\}$
$\alpha = \frac{\sqrt{\pi}}{4mrt}$
$\sigma = p\alpha$

Table 1. Conditions assumed for the parameters of the $PVW_{\mathcal{D}}$ scheme in our reduction.

we slightly modify the notation used in Algorithm 1 as described in Algorithm 3.

Theorem 4 (Decision LWE $\Longrightarrow$ FFP-NG). *Let $n, t, r, \ell, m, p, \alpha, \sigma$ as stated in Table 1. Let $\mathcal{A}$ be an FFP-NG adversary against the $PVW_{\mathcal{D}}$ encryption*

Algorithm 3 LWE-based PKE scheme for $\ell = 1$

- **Parameters:** Integers n, m, t, r, p and a real number $\alpha > 0$.
- **Private Key:** Choose $\mathbf{s} \in \mathbb{Z}_p^n$ uniformly at random. The private key is $\mathbf{s}$.
- **Public Key:** Choose $\mathbf{A} \in \mathbb{Z}_p^{n \times m}$ uniformly at random and $\mathbf{e} \in \mathbb{Z}_p^m$ by choosing each entry according to $\mathcal{D}_{\mathbb{Z}_p,\sigma}$. Set $\mathbf{b} = \mathbf{A}^T\mathbf{s} + \mathbf{e}$. The public key is $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_p^{n \times m} \times \mathbb{Z}_p^m$.
- **Encryption:** Given an element of the message space $v \in \mathbb{Z}_t$ and a public key $(\mathbf{A}, \mathbf{b})$, choose a vector $\mathbf{d} \in \{-r, \ldots, r\}^m$ uniformly at random. Set $\mathbf{u} = \mathbf{A}\mathbf{d}$ and $c = \langle \mathbf{b}, \mathbf{d} \rangle + \lfloor (p/t)\, \mathbf{v} \rceil$. The ciphertext is $(\mathbf{u}, c) \in \mathbb{Z}_p^n \times \mathbb{Z}_p$.
- **Decryption:** Given a ciphertext $(\mathbf{u}, c) \in \mathbb{Z}_p^n \times \mathbb{Z}_p$ and a private key $\mathbf{s} \in \mathbb{Z}_p^n$, output $\lfloor (t/p)\,(c - \langle \mathbf{s}, \mathbf{u} \rangle) \rceil$.

scheme, with error distribution $\mathcal{D}_{\mathbb{Z}_p,\sigma}$, denoted by Π. If the FFP-NG advantage $Adv_{\Pi}^{FFP-NG}(\mathcal{A}) \geq \delta$ is non-negligible in n, there exists an adversary $\mathcal{B}$ that solves the Decision LWE problem with error distribution $\mathcal{D}_{\mathbb{Z}_p,\hat{\sigma}}$, for $\hat{\sigma} = \sigma/\varphi$ such that

$$Adv_{\Pi}^{FFP\text{-}NG}(\mathcal{A}) \leq Adv_{\Pi}^{LWE}(\mathcal{B}) + \Gamma(n, \varphi), \tag{3}$$

where $\Gamma(n, \varphi) \in negl(n)$.

We provide an explicit bound for $\Gamma(n, \varphi)$ in Appendix A, illustrating the trade-off between the loss in LWE modulus-to-noise ratio, φ, and the advantage loss, of the reduction.

Before proving Theorem 4, we need three main ingredients. In Lemma 3 we determine a sufficient condition to get a decryption failure. In Lemma 4 we show that the reduction modulo p of a specific discrete Gaussian random variable does not affect the discrete value with overwhelming probability. In Lemma 5, under the assumption that the randomness $\mathbf{d}$ is independent of the error $\mathbf{e}$, we bound the tail of the distribution of $\left|\langle \mathbf{e}, \mathbf{d} \rangle\right|$.

The following lemma describes a sufficient condition and a necessary condition to get a decryption failure.

Lemma 3 (Decryption failure conditions). *Let $v \in \mathbb{Z}_t$ be a message, $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma}$, $\mathbf{d} \leftarrow \{-r, \ldots, r\}^m$ and w be the integer that satisfies $\frac{1}{2^{w+1}} \leq \frac{t}{p} < \frac{1}{2^w}$. The two following facts are true*

1. *If $\left|\langle \mathbf{e}, \mathbf{d} \rangle\right| > \frac{p}{2t}\left(1 + \frac{1}{2^w}\right) \implies$ a decryption failure occurs.*
2. *If a decryption occurs $\implies \left|\langle \mathbf{e}, \mathbf{d} \rangle\right| > \frac{p}{2t}\left(1 - \frac{1}{2^w}\right)$.*

then a decryption failure occurs.

Proof. Part 1. We use the assumption on w to prove that

$$\left| \frac{t}{p} \left\lfloor \frac{p}{t} v \right\rceil - v \right| < \frac{1}{2^{w+1}} \ .$$

Indeed, we have

$$\left|\frac{t}{p}\left\lfloor\frac{p}{t}v\right\rceil - v\right| = \frac{t}{p}\left|\left\lfloor\frac{p}{t}v\right\rceil - \frac{p}{t}v\right| < \frac{1}{2^w}\left|\left\lfloor\frac{p}{t}v\right\rceil - \frac{p}{t}v\right| \leq \frac{1}{2^{w+1}} \ .$$

The first inequality follows from the definition of w and last inequality follows directly from the definition of $\lfloor\cdot\rceil$. Now we prove the statement. We have

$$\left|\frac{t}{p}\langle\mathbf{e},\mathbf{d}\rangle + \frac{t}{p}\left\lfloor\frac{p}{t}v\right\rceil - v\right| \geq \left|\left|\frac{t}{p}\langle\mathbf{e},\mathbf{d}\rangle\right| - \left|\frac{t}{p}\left\lfloor\frac{p}{t}v\right\rceil - v\right|\right| > \left|\frac{1}{2} + \frac{1}{2^{w+1}} - \frac{1}{2^{w+1}}\right| = \frac{1}{2}$$

where the first inequality is the reverse triangle inequality, while in the second we use both the hypothesis and the inequality proved above. This means that a decryption failure occurs.

Part 2. We assume that a decryption failure occurs. This means that, given a message $v \in \mathbb{Z}_t$, we have

$$\left|\frac{t}{p}\langle\mathbf{e},\mathbf{d}\rangle + \frac{t}{p}\left\lfloor\frac{p}{t}v\right\rceil - v\right| \geq \frac{1}{2}.$$

By using the triangle inequality and the definition of rounding we have

$$\left|\frac{t}{p}\langle\mathbf{e},\mathbf{d}\rangle + \frac{t}{p}\left\lfloor\frac{p}{t}v\right\rceil - v\right| \leq \frac{t}{p}\left|\langle\mathbf{e},\mathbf{d}\rangle\right| + \frac{t}{p}\left|\left\lfloor\frac{p}{t}v\right\rceil - \frac{p}{t}v\right| \leq \frac{t}{p}\left|\langle\mathbf{e},\mathbf{d}\rangle\right| + \frac{t}{2p}.$$

By putting together the inequalities and by using the definition of w we have

$$\frac{t}{p}\left|\langle\mathbf{e},\mathbf{d}\rangle\right| + \frac{t}{2p} > \frac{1}{2} \iff \left|\langle\mathbf{e},\mathbf{d}\rangle\right| > \frac{p}{2t}\left(1 - \frac{t}{p}\right) \implies \left|\langle\mathbf{e},\mathbf{d}\rangle\right| > \frac{p}{2t}\left(1 - \frac{1}{2^w}\right).$$

□

Now, we prove that we can ignore the reduction modulo p. Assume the reduction modulo p returns elements in $\{-\frac{p-1}{2}, \ldots, \frac{p-1}{2}\}$. In this case, if we can prove that a random variable assumes values outside $\left[-\frac{p-1}{2}, \ldots, \frac{p-1}{2}\right]$ with negligible probability, we can ignore the effect of the reduction modulo p.

Lemma 4 (Ignoring reduction modulo p). *Let $\mathcal{X}$ be a random variable. If $\mathcal{X} \sim \mathcal{D}_{\mathbb{Z},\sigma}$, then*

$$\Pr\left[|\mathcal{X}| > \frac{p-1}{2}\right] \leq \frac{p-1}{2\sigma}\sqrt{2\pi\mathrm{e}} \cdot \mathrm{e}^{-\pi\left(\frac{p-1}{2\sigma}\right)^2} \ , \tag{4}$$

that is negligible in n.

Proof. We want to use the tail bound given in Equation 1, thus we must check that $(p-1)/2\sigma > 1/\sqrt{2\pi}$. By using the definition of σ and α

$$\begin{aligned} &\frac{p-1}{2\sigma} > \frac{1}{\sqrt{2\pi}} \\ \iff &p - 1 > \frac{2p\alpha}{\sqrt{2\pi}} \\ \iff &p - 1 > \frac{p}{2\sqrt{2}mrt}. \end{aligned}$$

Since $p > 2$ and $mrt \geq 2n^c$, we get

$$p - 1 > \frac{p}{4\sqrt{2}n^c} \geq \frac{p}{2\sqrt{2}mrt},$$

Now, we observe that for any positive number $s > 0$ and lattice Λ, we can define the lattice $\Lambda' = \Lambda/s$. Thus, for every x

$$\Pr\left[\mathcal{D}_{\Lambda,s} = sx\right] = \frac{\rho_s(sx)}{\rho_s(\Lambda)} = \frac{\rho(x)}{\rho(\Lambda')} = \Pr\left[\mathcal{D}_{\Lambda'} = x\right] \ .$$

Then, by taking $\mathcal{X}' \sim \mathcal{D}_{\mathbb{Z}/\sigma}$, we get

$$\Pr\left[\left|\mathcal{X}\right| > \frac{p-1}{2}\right] = \Pr\left[\left|\mathcal{X}'\right| > \frac{p-1}{2\sigma}\right].$$

We can apply Equation 1 and get

$$\Pr\left[\left|\mathcal{X}'\right| > \frac{p-1}{2\sigma}\right] < C \text{ , where } \quad C = \frac{p-1}{2\sigma}\sqrt{2\pi \mathrm{e}} \cdot \mathrm{e}^{-\pi\left(\frac{p-1}{2\sigma}\right)^2}.$$

We only need to check that $\frac{p-1}{2\sigma} \in \omega(\ln(n))$ (throughout the paper, when we check a limit of a function of n, we want the function to diverge faster than $\ln(n)$. In this way we can guarantee that our bounds are negligible in n). This is easily done by noticing that

$$\begin{aligned}\lim_{n\to\infty} \frac{p-1}{2\sigma \ln(n)} &= \lim_{n\to\infty} \frac{4(p-1)}{2p} \cdot \frac{mrt}{\ln(n)\sqrt{\pi}} \\ &\geq \lim_{n\to\infty} \frac{mrt}{\ln(n)\sqrt{\pi}},\end{aligned}$$

where for the inequality we have observed that $p \geq n^{c'}$ for some constant c' and thus $p \geq 2$ for sufficiently large n. Since r, t and φ are constants and $m = n^c$ with $c > 1$, it follows that

$$\lim_{n\to\infty} \frac{p-1}{2\sigma \ln(n)} = \infty.$$

Thanks to this, we get

$$\frac{p-1}{2\sigma}\sqrt{2\pi \mathrm{e}} \cdot \mathrm{e}^{-\pi\left(\frac{p-1}{2\sigma}\right)^2} \in negl(n).$$

□

The previous lemma guarantees that we can ignore the reduction modulo p. Indeed, with overwhelming probability, it will not affect the result. Looking ahead, we can use this to allow sampling from $\mathbb{Z}^m$ and viewing the result as an element of $\mathbb{Z}_p^m$ in the proof of Theorem 4, bearing in mind that the difference is negligible.

We want to show the only strategy $\mathcal{A}$ can use to win the FFP-NG game is to return a message-randomness pair $(v, \mathbf{d})$ for which the probability of obtaining a decryption failure is higher with respect to the public key $\mathcal{A}$ sees than with respect to an independent public key. To do so, we first show that the probability of getting a decryption failure for an independent public key is negligible in n. Let w be the integer such that $\frac{1}{2^{w+1}} \leq \frac{t}{p} < \frac{1}{2^w}$. Recall the sufficient condition for a decryption failure from Lemma 3.

Lemma 5 (Randomness independent of the error). *Given an error* $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma}$*, and a randomness* $\mathbf{d} \leftarrow \{-r, \ldots, r\}^m$*. If the randomness* $\mathbf{d}$ *is independent of the error* $\mathbf{e}$*, then*

$$\Pr_{\substack{\mathbf{e}\leftarrow\mathcal{D}_{\mathbb{Z}^m,\sigma}\\ \mathbf{d}\sim U}}\left[|\langle \mathbf{e}, \mathbf{d}\rangle| \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right)\right] \in negl\,(n)\,.$$

Proof. During this proof we write $\Pr[\cdot]$, implicitly saying that the probability is taken over $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma}$ and $\mathbf{d} \sim U$. Since $\mathbf{e}$ and $\mathbf{d}$ are independent, we want to apply Theorem 3 to show that $\langle \mathbf{e}, \mathbf{d}\rangle$ is ε-distributed according to $\mathcal{D}_{\mathbb{Z},\sigma\|\mathbf{d}\|}$, for some ε negligible in n. By using the same notation of Theorem 3, we set $\mathbf{z} = \mathbf{d}$, for $i = 1, \ldots, m$ we define $A_i = \Lambda_i = \mathbb{Z}$ and $s_i = \sigma$. Since all random variables are sampled from the same probability distribution, we know that $\Lambda_\cap = \Lambda_1 \cap \ldots \cap \Lambda_m = \mathbb{Z}$ is full rank. We define j to be the index that minimizes $|d_j| \neq 0$. To apply the theorem we must show

$$\left(\left(\frac{d_j}{\sigma}\right)^2 + \max_{i\neq j}\left(\frac{d_i}{\sigma}\right)^2\right)^{-\frac{1}{2}} \geq \eta_\varepsilon\,(\mathbb{Z})\,. \tag{5}$$

Using Theorem 2, we show below that

$$\left(\left(\frac{d_j}{\sigma}\right)^2 + \max_{i\neq j}\left(\frac{d_i}{\sigma}\right)^2\right)^{-\frac{1}{2}} \geq \sqrt{\frac{\ln\left(2\left(1 + 1/\varepsilon\right)\right)}{\pi}}\,.$$

Since $d \leftarrow \{-r, \ldots, r\}^m$, we can write

$$\begin{aligned}
\left(\left(\frac{d_j}{\sigma}\right)^2 + \max_{i\neq j}\left(\frac{d_i}{\sigma}\right)^2\right)^{-\frac{1}{2}} &= \left((\sigma)^{-2}\left(d_j^2 + \max_{i\neq j} d_i^2\right)\right)^{-\frac{1}{2}} \\
&= \sigma\left(d_j^2 + \max_{i\neq j} d_i^2\right)^{-\frac{1}{2}} \\
&\geq \frac{p\alpha}{\sqrt{2}r} \\
&= \frac{p}{4mr^2t}\sqrt{\frac{\pi}{2}} \\
&\geq \frac{\ln(n)}{2r}\sqrt{\frac{\pi}{2}}\,.
\end{aligned}$$

In the first inequality we exploit that $d_i,\ d_j < r$ and the definition of σ. The last inequality follows from the assumption that $p \geq 2mtr\ln(n)$. We see if the inequality is satisfied for value of ε that are negligible in n

$$\frac{\ln(n)}{2r}\sqrt{\frac{\pi}{2}} \geq \sqrt{\frac{\ln\left(2\left(1+1/\varepsilon\right)\right)}{\pi}} \iff \varepsilon \geq \frac{2}{\mathrm{e}^{\frac{\pi^2}{8r^2}\ln(n)^2} - 2}$$

that is negligible in n.

Now, we apply Theorem 3 to say that $\langle \mathbf{e}, \mathbf{d}\rangle \overset{\varepsilon'}{\approx} \mathcal{D}_{\mathbb{Z},\sigma\|\mathbf{d}\|}$. We have

$$\Pr\left[|\langle \mathbf{e}, \mathbf{d}\rangle| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)\right] \leq (1+\varepsilon')\cdot \Pr\left[|Z| \geq \frac{p}{2t\sigma\|\mathbf{d}\|}\left(1-\frac{1}{2^w}\right)\right]$$

where $Z \leftarrow \mathcal{D}_{\mathbb{Z}/(\sigma\|\mathbf{d}\|)}$. To apply Lemma 1 we must show

$$\frac{p}{2t\sigma\|\mathbf{d}\|}\left(1-\frac{1}{2^w}\right) \geq \frac{1}{\sqrt{2\pi}}\ .$$

To get a useful bound, we also have to check that

$$\frac{p}{2t\sigma\|\mathbf{d}\|}\left(1-\frac{1}{2^w}\right) \in \omega\left(\ln(n)\right)\ .$$

By using the definition of σ and α, we can write

$$\lim_{n\to\infty} \frac{p}{2t\sigma\|\mathbf{d}\|\ln(n)}\left(1-\frac{1}{2^w}\right) = \lim_{n\to\infty} \frac{2mr}{\ln(n)\|\mathbf{d}\|\sqrt{\pi}}\cdot\left(1-\frac{1}{2^w}\right)\ . \tag{6}$$

Since $\|\mathbf{d}\| \leq r\sqrt{m}$ and $m = n^c$ for $c > 1$, the first factor goes to infinity, while the term in brackets is in between $1/2$ and 1. Finally, we apply Lemma 1 and get

$$\begin{aligned}
\Pr\left[|\langle \mathbf{e}, \mathbf{d}\rangle| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)\right] &< (1+\varepsilon')\frac{p\sqrt{2\pi\mathrm{e}}}{2t\sigma\|\mathbf{d}\|}\left(1-\frac{1}{2^w}\right)\cdot \mathrm{e}^{-\pi\left(\frac{p}{2t\sigma\|\mathbf{d}\|}\left(1-\frac{1}{2^w}\right)\right)^2} \\
&= (1+\varepsilon')\frac{2\sqrt{2\mathrm{e}}mr}{\|\mathbf{d}\|}\left(1-\frac{1}{2^w}\right)\cdot \mathrm{e}^{-\pi\left(\frac{2mr}{\|\mathbf{d}\|\sqrt{\pi}}\left(1-\frac{1}{2^w}\right)\right)^2} \\
&< (1+\varepsilon')\frac{2\sqrt{2\mathrm{e}}mr}{\|\mathbf{d}\|}\cdot \mathrm{e}^{-\pi\left(\frac{2mr}{\|\mathbf{d}\|\sqrt{\pi}}\left(1-\frac{1}{2^w}\right)\right)^2}
\end{aligned}$$

that is negligible in terms of n. Indeed

$$\frac{2\sqrt{2\mathrm{e}}mr}{\|\mathbf{d}\|} = 2\sqrt{2\mathrm{e}}r\frac{n^c}{\|\mathbf{d}\|} < 2\sqrt{2\mathrm{e}}rn^c.$$

Finally,

$$\mathrm{e}^{-\pi\left(\frac{2mr}{\|\mathbf{d}\|\sqrt{\pi}}\left(1-\frac{1}{2^w}\right)\right)^2} \in \mathrm{e}^{-\omega\left((\ln n)^2\right)}$$

thanks to Equation (6).

□

Now, we analyze the winning probability of an FFP-NG adversary $\mathcal{A}$ against the LWE-based PKE Π. We describe an equivalent version of the FFP-NG game, called FFP-NG$'$, where $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$. We describe the details in Algorithm 4. We make explicit the intermediate state st between the algorithms $\mathcal{A}_0$ and $\mathcal{A}_1$.

Algorithm 4 Game FFP-NG$'$ equivalent to FFP-NG game against a PKE Π

$(sk_0, pk_0) \leftarrow \text{KeyGen}$
$(sk_1, pk_1) \leftarrow \text{KeyGen}$
$(v, \mathbf{d}, st) \leftarrow \mathcal{A}_0\,(pk_0)$
$b \leftarrow \{0,1\}$
$F_0 = \text{FCO}_0\,(v, \mathbf{d})$
$F_1 = \text{FCO}_1\,(v, \mathbf{d})$
$b' \leftarrow \mathcal{A}_1\,(st, F_b)$
return $[\![b = b']\!]$

By using this game we write

$$\Pr\left[\text{FFP-NG}_{\mathcal{A},\Pi} = 1\right] = \Pr\left[\text{FFP-NG}'_{\mathcal{A},\Pi} = 1\right] = \frac{1}{2}\sum \Pr\left[\text{FFP-NG}'_{\mathcal{A},\Pi} = 1 \middle| (b, F_0, F_1)\right] \Pr\left[(F_0, F_1)\,|b\right]$$

where the summation is over $(b, F_0, F_1) \in \{0,1\}^3$. If $F_0 = F_1$ we have

$$\Pr\left[0 \leftarrow \mathcal{A}_1\,(st, F_b) \,|\, (b, F_0, F_1)\right] = \Pr\left[0 \leftarrow \mathcal{A}_1\,(st, F_b) \,|\, (F_0, F_1)\right]\ ,$$
$$\Pr\left[1 \leftarrow \mathcal{A}_1\,(st, F_b) \,|\, (b, F_0, F_1)\right] = \Pr\left[1 \leftarrow \mathcal{A}_1\,(st, F_b) \,|\, (F_0, F_1)\right]\ .$$

In this case we have

$$\begin{aligned}
&\Pr\left[\text{FFP-NG}'_{\mathcal{A},\Pi} = 1 \middle| (F_0, F_1)\right] \\
&= \frac{1}{2}\left(\Pr\left[0 \leftarrow \mathcal{A}_1\,(st, F_b) \,|\, (b = 0, F_0, F_1)\right] + \Pr\left[1 \leftarrow \mathcal{A}_1\,(st, F_b) \,|\, (b = 1, F_0, F_1)\right]\right) \\
&= \frac{1}{2}\left(\Pr\left[0 \leftarrow \mathcal{A}_1\,(st, F_b) \,|\, (b = 0, F_0, F_1)\right] + 1 - \Pr\left[0 \leftarrow \mathcal{A}_1\,(st, F_b) \,|\, (b = 1, F_0, F_1)\right]\right) \\
&= \frac{1}{2}\ .
\end{aligned}$$

Thus, if $F_0 = F_1$ the adversary can only guess. Furthermore, thanks to Lemma 5 we know that $\Pr\left[F_1 = 1\right]$ is negligible in n. This implies that

$$\Pr\left[\text{FFP-NG}'_{\mathcal{A},\Pi} = 1 \middle| (F_0, F_1)\right] \Pr\left[F_0 = 0, F_1 = 1\right] \leq \Pr\left[F_1 = 1\right]$$

is negligible in n. We have just shown that the only possible strategy the adversary $\mathcal{A}$ can carry out to have a non-negligible advantage is to try to output a message-randomness pair that triggers a decryption failure with respect to the public key $\mathcal{A}$ sees. From now on, we can assume this without loss of generality. We are ready to define the adversary $\mathcal{B}$ against the Decision LWE and prove Theorem 4.

Proof (of Theorem 4). By using the decryption failure condition described in Lemma 3, we define the adversary $\mathcal{B}$ as described in Algorithm 5.

Algorithm 5 Adversary against the Decision LWE problem

Input: the public key $pk = (\mathbf{A}, \mathbf{b})$, where $\mathbf{A} \in \mathbb{Z}_p^{n \times m}$ and $\mathbf{b} \in \mathbb{Z}_p^m$.
Output: $b \leftarrow \{0, 1\}$.

Let $\varphi > 1$
$\mathbf{e}' \leftarrow \mathcal{D}_{\mathbb{Z}_p^m, \sigma'}$, where $\sigma' = \sigma\sqrt{1 - \frac{1}{\varphi^2}}$
$pk' \leftarrow (\mathbf{A}, \mathbf{b} + \mathbf{e}')$
$(v, \mathbf{d}) \leftarrow \mathcal{A}(pk')$
$k = \frac{p\|\mathbf{d}\|}{4rt}\sqrt{\frac{5}{6m}}$
if $|\langle \mathbf{e}', \mathbf{d} \rangle| \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right) - k$ **then**
    return 1 ▷ LWE Distribution
else
    return 0 ▷ Uniform Distribution
end if

To proving that $\mathcal{B}$ solves the decision LWE problem we need to show that

$$\left| \Pr_{(\mathbf{A},\mathbf{b}) \leftarrow A_{s, \mathcal{D}_{\mathbb{Z}_p^m}, \hat{\sigma}}} [\mathcal{B}(\mathbf{A}, \mathbf{b}) = 1] - \Pr_{(\mathbf{A},\mathbf{b}) \leftarrow U} [\mathcal{B}(\mathbf{A}, \mathbf{b}) = 1] \right|$$

is non-negligible in n. By the definition of $\mathcal{B}$, this is equivalent to prove that

$$\left| \Pr_{(\mathbf{A},\mathbf{b}) \leftarrow A_{s, \mathcal{D}_{\mathbb{Z}_p^m}, \hat{\sigma}}} \left[\left| \langle \mathbf{e}', \mathbf{d} \rangle \right| \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right) - k \right] - \Pr_{(\mathbf{A},\mathbf{b}) \leftarrow U} \left[\left| \langle \mathbf{e}', \mathbf{d} \rangle \right| \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right) - k \right] \right| \tag{7}$$

is non-negligible in n. We analyze the two terms separately. In Lemma 6 we show that the second term of Equation (7) is negligible in n, while in Lemma 7 we show that the first term of Equation (7) is non-negligible in n. Putting together the two results, we prove that the absolute difference in Equation (7) is non-negligible in n. In Appendix A we also provide the analysis of the negligible function $\Gamma(n, \varphi)$ and how to get the inequality described in Equation 3.

□

We first study the second term of Equation (7) and show the following result.

Lemma 6. *Let $\mathbf{e}' \sim \mathcal{D}_{\mathbb{Z}^m, \sigma}$ be an error, and $\mathbf{d} \leftarrow \{-r, \ldots, r\}^m$ be a randomness. If $(\mathbf{A}, \mathbf{b})$ is sampled uniformly at random, then*

$$\Pr_{(\mathbf{A},\mathbf{b}) \leftarrow U} \left[\left| \langle \mathbf{e}', \mathbf{d} \rangle \right| \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right) - k \right] \in negl(n). \tag{8}$$

Proof. The proof of Lemma 6 is the same as the proof of Lemma 5. Indeed, being $(\mathbf{A}, \mathbf{b})$ chosen uniformly at random, we get that $\mathbf{e}'$ and $\mathbf{d}$ are independent. The only other two differences are that we use a slightly different probability distribution and now our tail is shorter. Instead of sampling from $\mathcal{D}_{\mathbb{Z}_p^m,\sigma}$, we sample $\mathbf{e}'$ from $\mathcal{D}_{\mathbb{Z}_p^m,\sigma'}$ and, instead of considering values $\geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)$, we consider values $\geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right) - k$.

Now we analyze the first term of Equation (7) and prove the following result.

Lemma 7. *Let $\mathbf{e}' \sim \mathcal{D}_{\mathbb{Z}^m,\sigma}$ be an error, and $\mathbf{d} \leftarrow \{-r,\ldots,r\}^m$ be a randomness. If $(\mathbf{A},\mathbf{b})$ is sampled according to $A_{s,\mathcal{D}_{\mathbb{Z}_p^m,\hat{\sigma}}}$, then*

$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow A_{s,\mathcal{D}_{\mathbb{Z}_p^m,\hat{\sigma}}}}\left[\left|\langle \mathbf{e}', \mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right) - k\right] \notin negl\left(n\right). \tag{9}$$

Proof. From now on, unless explicitly written, we will denote $\Pr_{(\mathbf{A},\mathbf{b})\leftarrow A_{s,\mathcal{D}_{\mathbb{Z}_p^m,\hat{\sigma}}}}[\cdot]$ by $\Pr[\cdot]$. Let $\tilde{b}$ be the bit in the FFP-NG$'$ game. Thus, the adversary $\mathcal{A}$ gets the public key $pk_{\tilde{b}}$ in the FFP-NG$^{\mathcal{A}}_{\Pi}$ game. We have

$$\begin{aligned}
&\Pr\left[\left|\langle \mathbf{e}', \mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right) - k\right] = \\
&\Pr\left[\left|\langle \mathbf{e}', \mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right) - k \,\middle|\, \text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1 \wedge \tilde{b} = 0\right]\Pr\left[\text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1 \wedge \tilde{b} = 0\right] + \\
&\Pr\left[\left|\langle \mathbf{e}', \mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right) - k \,\middle|\, \text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1 \wedge \tilde{b} = 1\right]\Pr\left[\text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1 \wedge \tilde{b} = 1\right] + \\
&\Pr\left[\left|\langle \mathbf{e}', \mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right) - k \,\middle|\, \text{FFP-NG}^{\mathcal{A}}_{\Pi} = 0\right]\Pr\left[\text{FFP-NG}^{\mathcal{A}}_{\Pi} = 0\right] \geq \\
&\Pr\left[\left|\langle \mathbf{e}', \mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right) - k \,\middle|\, \text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1 \wedge \tilde{b} = 0\right]\Pr\left[\text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1 \wedge \tilde{b} = 0\right].
\end{aligned}$$

We define the error $\mathbf{e} = \mathbf{e}' + \hat{\mathbf{e}}$, this is the error in the public key the adversary $\mathcal{A}$ receives as input. Note that, by the assumption that $\mathcal{A}$ wants to obtain a decryption failure for the key it sees and since we are conditioning on $\mathcal{A}$ winning the game, we have that

$$\text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1 \wedge \tilde{b} = 0 \Longrightarrow \left|\langle \mathbf{e}, \mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right).$$

We can use this observation and the fact that $\mathbf{e} = \mathbf{e}' + \hat{\mathbf{e}}$ to say

$$\begin{aligned}
&\Pr\left[\left|\langle \mathbf{e}', \mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right) - k \,\middle|\, \text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1 \wedge \tilde{b} = 0\right] \geq \\
&\Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d}\rangle\right| \leq k \,\middle|\, \text{FFP-NG}^{\mathcal{A}}_{\Pi} = 1 \wedge \tilde{b} = 0\right].
\end{aligned}$$

Instead of proving Equation (9), we can prove that the latter probability is non-negligible in n or, equivalently, we can show that

$$\Pr\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k \;\middle|\; \text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right] \in negl\,(n)\,.$$

By assumption, the FFP-NG advantage of $\mathcal{A}$ is greater then a non-negligible value, namely δ. We can show that $\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right] \geq \delta$ as well. Indeed, we have

$$\begin{aligned}
\frac{1}{2} + \delta &\leq \frac{1}{2} + \text{Adv}_{\mathcal{A},\Pi}^{\text{FFP-NG}} \\
&= \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1\right] \\
&= \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 1\right] + \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right] \\
&= \frac{1}{2}\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \;\middle|\; \tilde{b} = 1\right] + \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right] \\
&= \frac{1}{2}\left(1 - \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 0 \;\middle|\; \tilde{b} = 1\right]\right) + \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right] \\
&\leq \frac{1}{2} + \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right]. \qquad (10)
\end{aligned}$$

For the first inequality we use the assumption on $\text{Adv}_{\mathcal{A},\Pi}^{\text{FFP-NG}}$. For the third equality we use

$$\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \;\wedge \tilde{b} = 1\right] = \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \;\middle|\; \tilde{b} = 1\right]\Pr\left[\tilde{b} = 1\right]$$

and $\Pr\left[\tilde{b} = 1\right] = 1/2$. The last inequality is simply due to

$$\left(1 - \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 0 \;\middle|\; \tilde{b} = 1\right]\right) \leq 1.$$

By comparing the left hand side and the last right hand side of the chain of inequality above, we get exactly what we want. Now, if we show that

$$\Pr\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k\right] \in negl\,(n)$$

we get that

$$\begin{aligned}
&\Pr\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k \;\middle|\; \text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right] \\
=&\frac{\Pr\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k \wedge \text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right]}{\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right]} \\
\leq&\frac{1}{\delta}\Pr\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k \wedge \text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right] \\
\leq&\frac{1}{\delta}\Pr\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k\right]
\end{aligned}$$

is negligible in n as well. For the first inequality we have used Equation (10).

We go on to show that

$$\Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d}\rangle\right| > k\right] \in negl\,(n)\,. \tag{11}$$

Recall that $\mathbf{e} = \hat{\mathbf{e}} + \mathbf{e}'$, where $\hat{\mathbf{e}} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\hat{\sigma}}$ and $\mathbf{e}' \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma'}$. We want to prove that $\mathbf{e}$ is ε-distributed according to $\mathcal{D}_{\mathbb{Z}^m,\sqrt{(\hat{\sigma})^2+(\sigma')^2}}$, for some ε negligible in n. We thus analyze the probability distribution of $\hat{\mathbf{e}}$ conditioned on $\mathbf{e}$, since we know that $\hat{\mathbf{e}} - \mathbf{e} - \mathbf{d}$ is a Markov chain. Theorem 2 gives us both answers. Each coordinate of $\mathbf{e}'$ and $\hat{\mathbf{e}}$ are independent, thus we can set $\Sigma_1 = (\hat{\sigma})^2 \cdot I_m$, $\Sigma_2 = (\sigma')^2 \cdot I_m$, $A_1 = A_2 = \Lambda_1 = \Lambda_2 = \mathbb{Z}^m$. This implies $\Lambda_\cap = \mathbb{Z}^m$. We denote $A = A_1 + A_2 = \mathbb{Z}^m$ and $\Sigma = \sigma^2 \cdot I_m$, with $\sigma = \sqrt{(\hat{\sigma})^2 + (\sigma')^2}$. To apply the theorem we must prove that

$$\bar{\sigma} \geq \eta_\varepsilon\,(\mathbb{Z}^m) \qquad \text{and} \qquad \sigma' \geq \eta_\varepsilon\,(\mathbb{Z}^m)$$

where $\bar{\sigma} = \sigma\frac{\sqrt{\varphi^2-1}}{\varphi^2}$ and $\Sigma_3 = (\bar{\sigma})^2 \cdot I_m$.
By using again Theorem 2, we can show that

$$\bar{\sigma} \geq \sqrt{\frac{\ln\left(2m\left(1+1/\varepsilon\right)\right)}{\pi}} \qquad \text{and} \qquad \sigma' \geq \sqrt{\frac{\ln\left(2m\left(1+1/\varepsilon\right)\right)}{\pi}}.$$

We are going to show computations only for one of the two, the other is the same. By using the definition of σ' and the condition on p in Table 1 we show

$$\sigma' = \frac{p}{4\varphi mrt}\sqrt{\pi\left(\varphi^2-1\right)} \geq \frac{\ln(n)}{2\varphi}\sqrt{\pi\left(\varphi^2-1\right)}.$$

We verify the inequality

$$\begin{aligned}
\frac{\ln(n)}{2\varphi}\sqrt{\pi\left(\varphi^2-1\right)} \geq \sqrt{\frac{\ln\left(2m\left(1+1/\varepsilon\right)\right)}{\pi}} \iff \\
\frac{\ln(n)^2}{4\varphi^2}\pi\left(\varphi^2-1\right) \geq \frac{\ln\left(2m\left(1+1/\varepsilon\right)\right)}{\pi} \iff \\
\varepsilon \geq \frac{2m}{\mathrm{e}^{\frac{\pi^2\left(\varphi^2-1\right)}{4\varphi^2}\ln(n)^2} - 2m}.
\end{aligned}$$

In the first step, we just take the square of both sides of the inequality. In the second step, we exponentiate both sides of the inequality and isolate ε. The right hand side of the last inequality is negligible in n. The same is true for $\bar{\sigma}$. Thus, we can apply Theorem 2. By using the first part of the theorem we get

$$\chi_2 \overset{\varepsilon'}{\approx} \mathcal{D}_{\mathbb{Z}^m,\sqrt{(\sigma')^2+(\hat{\sigma})^2}} = \mathcal{D}_{\mathbb{Z}^m,\sigma},$$

and we just need to notice that $\mathbf{e} \leftarrow \chi_2$. Thus, we can say that $\mathbf{e}$ is ε'-distributed as $\mathcal{D}_{\mathbb{Z}^m,\sigma}$.
By using the second part of the theorem we get

$$\chi_1^{\mathbf{e}} \overset{\bar{\varepsilon}}{\approx} \frac{\bar{\mathbf{e}}}{\varphi^2} + \mathcal{D}_{\mathbb{Z}^m - \frac{\bar{\mathbf{e}}}{\varphi^2},\bar{\sigma}},$$

and notice that $\chi_1^{\mathbf{e}}$ is the probability distribution of $\hat{\mathbf{e}}$ conditioned on $\mathbf{e}$.
Now we want to set a threshold, θ, such that:

1. If $\|\bar{\mathbf{e}}\| < \theta \Longrightarrow \left|\langle \hat{\mathbf{e}}, \mathbf{d} \rangle\right| > k$ with negligible in n probability;
2. If $\|\mathbf{e}\| \geq \theta \Longrightarrow \mathbf{e} = \bar{\mathbf{e}}$ with negligible in n probability.

We set

$$\theta = \frac{p}{2rt}\sqrt{\frac{1}{6m}}.$$

Now we are ready to prove Equation (11). We can write

$$\begin{aligned}
\Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d} \rangle\right| > k\right] =& \Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d} \rangle\right| > k \middle| \|\mathbf{e}\| \geq \theta\right] \cdot \Pr\left[\|\mathbf{e}\| \geq \theta\right] \\
&+ \Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d} \rangle\right| > k \middle| \|\mathbf{e}\| < \theta\right] \cdot \Pr\left[\|\mathbf{e}\| < \theta\right] \\
\leq& \Pr\left[\|\mathbf{e}\| \geq \theta\right] + \Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d} \rangle\right| > k \middle| \|\mathbf{e}\| < \theta\right].
\end{aligned}$$

We now analyze both terms separately and show they are both negligible in n.
We start with $\Pr\left[\|\mathbf{e}\| \geq \theta\right]$. We can write

$$\Pr\left[\|\mathbf{e}\| \geq \theta\right] \leq (1+\varepsilon') \Pr\left[\|Z\| \geq \frac{\theta}{\sigma}\right] = (1+\varepsilon') \Pr\left[\|Z\| \geq \frac{\sqrt{m}\theta}{\sqrt{m}\sigma}\right]$$

with $Z \leftarrow \mathcal{D}_{\mathbb{Z}^m/\sigma}$. We want to use again Equation 1, to do so, we need to check that $\theta/\left(\sqrt{m}\sigma\right) \geq 1/\sqrt{2\pi}$. By using the definitions of θ and σ, we get

$$\begin{aligned}
\frac{\theta}{\sigma\sqrt{m}} =& \frac{4pmrt}{2\sqrt{6\pi}pmrt} \\
=& \sqrt{\frac{2}{3\pi}} \\
\geq& \frac{1}{\sqrt{2\pi}}.
\end{aligned}$$

So, $\Pr\left[\|\mathbf{e}\| \geq \theta\right] \in negl\left(n\right)$. Now we analyze the term

$$\Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d} \rangle\right| > k \middle| \|\mathbf{e}\| < \theta\right]$$

We denote with $B^m\left(0, \theta\right)$ the open m-dimensional ball centered in 0 and with radius θ. With this notation we can write

$$\Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d} \rangle\right| > k \middle| \|\mathbf{e}\| < \theta\right] = \sum_{\bar{\mathbf{e}} \in B^m(0,\theta)} \Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d} \rangle\right| > k \middle| \mathbf{e} = \bar{\mathbf{e}}\right] \Pr\left[\mathbf{e} = \bar{\mathbf{e}}\right].$$

We want to find a bound on $\Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d}\rangle\right| \geq k \middle| \mathbf{e} = \bar{\mathbf{e}}\right]$ that does not depend on $\bar{\mathbf{e}}$. We have

$$\begin{aligned}\Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d}\rangle\right| > k \middle| \mathbf{e} = \bar{\mathbf{e}}\right] &= \Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d}\rangle - \frac{1}{\varphi^2}\langle \bar{\mathbf{e}}, \mathbf{d}\rangle + \frac{1}{\varphi^2}\langle \bar{\mathbf{e}}, \mathbf{d}\rangle\right| > k \middle| \mathbf{e} = \bar{\mathbf{e}}\right] \\ &\leq \Pr\left[\left|\left\langle \hat{\mathbf{e}} - \frac{\bar{\mathbf{e}}}{\varphi^2}, \mathbf{d}\right\rangle\right| > k - \frac{1}{\varphi^2}\left|\langle \bar{\mathbf{e}}, \mathbf{d}\rangle\right| \middle| \mathbf{e} = \bar{\mathbf{e}}\right],\end{aligned}$$

where the inequality is the triangle inequality. Since we are conditioning on $\mathbf{e}$, we know that $\mathbf{d}$ is independent from $\hat{\mathbf{e}}$. We want to show that

$$\left\langle \hat{\mathbf{e}} - \frac{\bar{\mathbf{e}}}{\varphi^2}, \mathbf{d}\right\rangle \text{ is } \varepsilon'\text{-distributed according to } \left\langle \frac{\bar{\mathbf{e}}}{\varphi^2}, \mathbf{d}\right\rangle + D_{\mathbb{Z}-\langle \bar{\mathbf{e}}, \mathbf{d}\rangle/\varphi^2, \bar{\sigma}\|\mathbf{d}\|}$$

To do so, we want to use Theorem 3, so we have to check that

$$\left(\left(\frac{d_j}{\bar{\sigma}}\right)^2 + \max_{i \neq j}\left(\frac{d_i}{\bar{\sigma}}\right)^2\right)^{-\frac{1}{2}} \geq \sqrt{\frac{\ln\left(2\left(1 + 1/\varepsilon\right)\right)}{\pi}}.$$

Computations are similar to what we have done to prove Equation (5). We can now conclude by using for the last time the tail bound given in Equation (1)

$$\begin{aligned}&\Pr\left[\left|\langle \hat{\mathbf{e}}, \mathbf{d}\rangle\right| > k \middle| \mathbf{e} = \bar{\mathbf{e}}\right] \\ &\quad\leq (1 + \varepsilon')\Pr\left[\frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left|\left\langle \hat{\mathbf{e}} - \frac{\bar{\mathbf{e}}}{\varphi^2}, \mathbf{d}\right\rangle\right| > \frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k - \frac{1}{\varphi^2}\left|\langle \bar{\mathbf{e}}, \mathbf{d}\rangle\right|\right) \middle| \mathbf{e} = \bar{\mathbf{e}}\right] \\ &\quad\leq (1 + \varepsilon')\frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k - \frac{1}{\varphi^2}\left|\langle \bar{\mathbf{e}}, \mathbf{d}\rangle\right|\right) \cdot \sqrt{2\pi \mathrm{e}} \cdot \mathrm{e}^{-\pi\left(\frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k - \frac{1}{\varphi^2}\left|\langle \bar{\mathbf{e}}, \mathbf{d}\rangle\right|\right)\right)^2}.\end{aligned}$$

To apply the last inequality and to prove that it is negligible in n we have to check that

$$\frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k - \frac{1}{\varphi^2}\left|\langle \bar{\mathbf{e}}, \mathbf{d}\rangle\right|\right) \geq \frac{1}{\sqrt{2\pi}} \quad \text{and} \quad \frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k - \frac{1}{\varphi^2}\left|\langle \bar{\mathbf{e}}, \mathbf{d}\rangle\right|\right) \in \omega\left(\ln\left(n\right)\right) \ .$$

We can prove it both only rearranging the term. First of all, by using the definitions of k, θ and $\bar{\sigma}$, we get

$$\begin{aligned}\frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k - \frac{1}{\varphi^2}\left|\langle \bar{\mathbf{e}}, \mathbf{d}\rangle\right|\right) &\geq \frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k - \frac{\theta}{\varphi^2}\|\mathbf{d}\|\right) \\ &= \frac{1}{\sqrt{2\pi}}\left[2\varphi^2\sqrt{\frac{m}{3\left(\varphi^2 - 1\right)}}\left(\frac{\sqrt{5}}{2} - \frac{1}{\varphi^2}\right)\right]\end{aligned}$$

that belongs to $\omega\left(\ln\left(n\right)\right)$, since $m = n^c$, with $c > 1$. It is also easy to check that the value is greater than $1/\sqrt{2\pi}$. Indeed, we have

$$\left[2\varphi^2\sqrt{\frac{m}{3\left(\varphi^2 - 1\right)}}\left(\frac{\sqrt{5}}{2} - \frac{1}{\varphi^2}\right)\right] > 1$$

since each factor is > 1. This shows that Equation (11) holds and consequently proves the lemma.

□

The proof can be generalized easily to the case $1 \neq \ell \in \mathcal{O}(n)$. We switch again to the notation used in Algorithm 1.

Theorem 5. *Let $n, t, r, m, p, \alpha, \sigma$ as stated in Table 1 and $\ell \in \mathcal{O}(n)$. Let $\mathcal{A}$ be an FFP-NG adversary against the $PVW_{\mathcal{D}}$ encryption scheme, with error distribution $\mathcal{D}_{\mathbb{Z}_p,\sigma}$, denoted by Π. If the FFP-NG advantage $Adv_{\Pi}^{FFP-NG}(\mathcal{A}) \geq \delta$ is non-negligible in n, there exists an adversary $\mathcal{B}$ that solves the Decision LWE problem with error distribution $\mathcal{D}_{\mathbb{Z}_p,\hat{\sigma}}$, for $\hat{\sigma} = \sigma/\varphi$ such that*

$$Adv_{\Pi}^{FFP\text{-}NG}(\mathcal{A}) \leq Adv_{\Pi}^{LWE}(\mathcal{B}) + \bar{\Gamma}(n, \varphi),$$

where $\bar{\Gamma}(n, \varphi) \in negl(n)$.

The structure of the proof of this theorem is essentially the same of Theorem 4. We only highlight the main changes needed to prove this result. Since we work with ℓ dimensional messages, we must change the condition to get a decryption failure. We observe that both encryption and decryption are computed component wise. Hence, if there is an index j such that $|\langle \mathbf{e}_j, \mathbf{d} \rangle| \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right)$, we get a decryption failure on the j-th coordinate of the message. Thus, we can easily describe a sufficient condition for decryption failure. Given an error $\mathbf{E} \in \mathbb{Z}_p^{\ell \times m}$ and randomness $\mathbf{d} \in \{-r, \ldots, r\}^m$ we have

1. if $\|\mathbf{E}\mathbf{d}\|_\infty \geq \frac{p}{2t}\left(1 + \frac{1}{2^w}\right) \Longrightarrow$ decryption failure occurs.
2. If a decryption failure occurs $\Longrightarrow \|\mathbf{E}\mathbf{d}\|_\infty \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right)$.

To prove the equivalent of Lemma 6 and Lemma 7, we use the union bound and the symmetry between the rows of the error matrix as follows

$$\begin{aligned}\Pr\left[\|\mathbf{E}'\mathbf{d}\|_\infty \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right) - k\right] &\leq \sum_{i=1}^{\ell} \Pr\left[\left|\langle \mathbf{e}'_i, \mathbf{d} \rangle\right| \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right) - k\right] \\ &\leq \ell \cdot \Pr\left[\left|\langle \mathbf{e}'_1, \mathbf{d} \rangle\right| \geq \frac{p}{2t}\left(1 - \frac{1}{2^w}\right) - k\right]\end{aligned}$$

Observe that we reduce the problem to the $\ell = 1$ case analyzed in Theorem 4. Hence, the same argument can be used. The other change we need to do is the definition of the threshold θ used in Lemma 7. To do so, we need to introduce a matrix norm. For any matrix H we denote by $\|H\| = \max_i \|H^i\|_2$, where H^i is the i-th column of H. By using this norm, we can define the threshold θ that satisfies the following two conditions

1. If $\|\bar{\mathbf{E}}^T\| < \theta \Longrightarrow \left\|\hat{\mathbf{E}}\mathbf{d}\right\|_\infty \geq k$ with negligible in n probability;
2. If $\|\bar{\mathbf{E}}^T\| \geq \theta \Longrightarrow \mathbf{E} = \bar{\mathbf{E}}$ with negligible in n probability.

Thanks to these modifications and the union bound, we can repeat the proof of Theorem 4 to also prove Theorem 5. The analysis of $\bar{\Gamma}(n, \varphi)$ is similar the one in Appendix A.

Acknowledgement

FS and CM is 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).

A Advantages in Theorem 4

In this section we provide all details needed to compute the advantage introduced in Theorem 4. To lighten the notation we write $\Pr_{(\mathbf{A},\mathbf{b})\leftarrow \mathrm{LWE}}[\cdot]$ instead of $\Pr_{(\mathbf{A},\mathbf{b})\leftarrow A_{s,\mathcal{D}_{\mathbb{Z}_p^m},\hat{\sigma}}}[\cdot]$. We collect all computations that are involved in the analysis of the advantage of the LWE adversary $\mathcal{B}$ and highlight the relation with the advantage of the FFP-NG adversary $\mathcal{A}$. Since in our reduction we are ignoring the reduction modulo p the first loss comes from this choice. As stated in Lemma 4 the probability that we need the rounding, in this case negligible in n, can be bounded as follows

$$\begin{aligned}\Pr\left[|\mathcal{X}| > \frac{p-1}{2}\right] &\leq \frac{p-1}{2\sigma}\sqrt{2\pi \mathrm{e}}\cdot \mathrm{e}^{-\pi\left(\frac{p-1}{2\sigma}\right)^2}\\ &\leq 5mrt\cdot \mathrm{e}^{-\pi\left(\frac{p-1}{2\sigma}\right)^2}\\ &\leq 5mrt\cdot \mathrm{e}^{-(mrt)^2}.\end{aligned}$$

where $\mathcal{X}\sim\mathcal{D}_{\mathbb{Z},\sigma}$. For the inequalities we have used the definition of σ, $p>2$, and $2\sqrt{2\mathrm{e}}<5$. Since we deal with random variable over $\mathbb{Z}^m$, we can use the union bound to get the loss of ignoring the rounding as

$$\leq 5m^2rt\cdot \mathrm{e}^{-(mrt)^2},$$

that is negligible in n. Ignoring the reduction modulo p, we can focus on the advantage of the adversary $\mathcal{B}$. By definition of advantage we have

$$\mathrm{Adv}_{\Pi}^{\mathrm{LWE}}(\mathcal{B}) = \left|\Pr_{(\mathbf{A},\mathbf{b})\leftarrow \mathrm{LWE}}[\mathcal{B}(\mathbf{A},\mathbf{b})=1] - \Pr_{(\mathbf{A},\mathbf{b})\leftarrow U}[\mathcal{B}(\mathbf{A},\mathbf{b})=1]\right|. \tag{12}$$

By using the definition of $\mathcal{B}$ we can rewrite the advantage as

$$\left|\Pr_{(\mathbf{A},\mathbf{b})\leftarrow \mathrm{LWE}}\left[\left|\langle \mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right] - \Pr_{(\mathbf{A},\mathbf{b})\leftarrow U}\left[\left|\langle \mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right]\right|.$$

By following some computations in Lemma 7, we write

$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right] =$$
$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k \,\middle|\, \text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right]\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right]+$$
$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k \,\middle|\, \text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=1\right]\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=1\right]+$$
$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k \,\middle|\, \text{FFP-NG}_{\Pi}^{\mathcal{A}}=0\right]\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=0\right] \geq$$
$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k \,\middle|\, \text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right]\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right].$$

We now use that $\mathbf{e}=\mathbf{e}'+\hat{\mathbf{e}}$ and that we are conditioning on $\mathcal{A}$ winning the FFP-NG game, we can write

$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k \,\middle|\, \text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right] \geq$$
$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| \leq k \,\middle|\, \text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right].$$

Recall that, by assumption, the advantage of $\mathcal{A}$ is non-negligible in n. Thanks to this we have

$$\begin{aligned}
\frac{1}{2}+\text{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}) &= \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\right]\\
&= \Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=1\right]+\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right]\\
&= \frac{1}{2}\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\,\middle|\,\tilde{b}=1\right]+\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right]\\
&= \frac{1}{2}\left(1-\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=0\,\middle|\,\tilde{b}=1\right]\right)+\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right]\\
&\leq \frac{1}{2}+\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right].
\end{aligned}$$

In turn, we have shown that

$$\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right] \geq \text{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}),$$

that is non-negligible by assumption. Thanks to this equation we can write

$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| \leq k \,\middle|\, \text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right]\Pr\left[\text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right] \geq$$
$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| \leq k \,\middle|\, \text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right]\text{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}) =$$
$$\left(1-\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\text{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k \,\middle|\, \text{FFP-NG}_{\Pi}^{\mathcal{A}}=1\wedge\tilde{b}=0\right]\right)\text{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}).$$

Now, we collect all we have done so far

$$\mathrm{Adv}_{\Pi}^{\mathrm{LWE}}(\mathcal{B}) =$$

$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\mathrm{LWE}}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right] - \Pr_{(\mathbf{A},\mathbf{b})\leftarrow U}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right] \geq$$

$$\mathrm{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}) - \Pr_{(\mathbf{A},\mathbf{b})\leftarrow\mathrm{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k \;\middle|\; \text{FFP-NG}_{\Pi}^{\mathcal{A}} = 1 \wedge \tilde{b} = 0\right] \mathrm{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}) -$$

$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow U}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right] \geq$$

$$\mathrm{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}) - \Pr_{(\mathbf{A},\mathbf{b})\leftarrow\mathrm{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k\right] - \Pr_{(\mathbf{A},\mathbf{b})\leftarrow U}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right].$$

We can write

$$\mathrm{Adv}_{\Pi}^{\text{FFP-NG}}(\mathcal{A}) \leq$$

$$\mathrm{Adv}_{\Pi}^{\mathrm{LWE}}(\mathcal{B}) + \Pr_{(\mathbf{A},\mathbf{b})\leftarrow\mathrm{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k\right] + \Pr_{(\mathbf{A},\mathbf{b})\leftarrow U}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right].$$

In Lemma 6 we have shown that

$$\begin{aligned}
&\Pr_{(\mathbf{A},\mathbf{b})\leftarrow U}\left[\left|\langle\mathbf{e}',\mathbf{d}\rangle\right| \geq \frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right] \\
\leq&(1+\varepsilon')\frac{\sqrt{2\pi\mathrm{e}}}{\sigma'\|\mathbf{d}\|}\left(\frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right)\cdot\mathrm{e}^{-\pi\left(\frac{1}{\sigma'\|\mathbf{d}\|}\left(\frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right)\right)^2} \\
\leq&\sqrt{\frac{\varphi^2}{\varphi^2-1}}\frac{10mr}{\|\mathbf{d}\|}\cdot\mathrm{e}^{-\pi\left(\frac{1}{\sigma'\|\mathbf{d}\|}\left(\frac{p}{2t}\left(1-\frac{1}{2^w}\right)-k\right)\right)^2} \\
\leq&\sqrt{\frac{\varphi^2}{\varphi^2-1}}\frac{10mr}{\|\mathbf{d}\|}\mathrm{e}^{-\frac{\varphi^2 m}{\varphi^2-1}}
\end{aligned}$$

that is negligible in n. For the second and third inequality we have used the definition of σ' and k, $w>1$, and $4<2\sqrt{2\mathrm{e}}<5$. Here and in the following we have, and repeatedly will, use $\varepsilon'\leq 1$. While, in Lemma 7 we have shown that

$$\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\mathrm{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k\right] \leq \Pr_{(\mathbf{A},\mathbf{b})\leftarrow\mathrm{LWE}}\left[\|\mathbf{e}\| \geq \theta\right] + \Pr_{(\mathbf{A},\mathbf{b})\leftarrow\mathrm{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k \;\middle|\; \|\mathbf{e}\| < \theta\right],$$

Analyzing the first term, we have

$$\begin{aligned}
\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\mathrm{LWE}}\left[\|\mathbf{e}\| \geq \theta\right] \leq&(1+\varepsilon')\left(\frac{\theta\sqrt{2\pi\mathrm{e}}}{\sqrt{m}\sigma}\cdot\mathrm{e}^{-\pi\left(\frac{\theta}{\sqrt{m}\sigma}\right)^2}\right)^m \\
=&(1+\varepsilon')\left(\frac{2}{\sqrt{3}}\cdot\mathrm{e}^{-\frac{1}{6}}\right)^m \\
\leq&2\left(\frac{2}{\sqrt{3}}\cdot\mathrm{e}^{-\frac{1}{6}}\right)^m
\end{aligned}$$

that is negligible in n, since $\frac{2}{\sqrt{3}} \cdot \mathrm{e}^{-\frac{1}{6}} < 1$. For the second one, we have

$$
\begin{aligned}
&\Pr_{(\mathbf{A},\mathbf{b})\leftarrow\mathrm{LWE}}\left[\left|\langle\hat{\mathbf{e}},\mathbf{d}\rangle\right| > k \middle| \|\mathbf{e}\| < \theta\right] \\
&\leq (1+\varepsilon')\,\frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k - \frac{1}{\varphi^2}\left|\langle\bar{\mathbf{e}},\mathbf{d}\rangle\right|\right)\cdot\sqrt{2\pi\mathrm{e}}\cdot\mathrm{e}^{-\pi\left(\frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k-\frac{1}{\varphi^2}\left|\langle\bar{\mathbf{e}},\mathbf{d}\rangle\right|\right)\right)^2} \\
&\leq (1+\varepsilon')\,9\varphi^2\sqrt{\frac{m}{\varphi^2-1}}\cdot\mathrm{e}^{-\pi\left(\frac{1}{\bar{\sigma}\|\mathbf{d}\|}\left(k-\frac{1}{\varphi^2}\left|\langle\bar{\mathbf{e}},\mathbf{d}\rangle\right|\right)\right)^2} \\
&\leq (1+\varepsilon')\,9\varphi^2\sqrt{\frac{m}{\varphi^2-1}}\cdot\mathrm{e}^{-\frac{\varphi^4 m}{400(\varphi^2-1)}} \\
&\leq 18\varphi^2\sqrt{\frac{m}{\varphi^2-1}}\cdot\mathrm{e}^{-\frac{\varphi^4 m}{400(\varphi^2-1)}}
\end{aligned}
$$

that is negligible in n. For the second and third inequality we have used the definition of $\bar{\sigma}$, θ and k, $\varepsilon\prime < 1$, $4\sqrt{(5\mathrm{e})/3} < 9$, and $\sqrt{5/6} - \sqrt{2/3} > 1/20$. By collecting the different bound we can define the function $\mathit{\Gamma}$ as follows

$$
\begin{aligned}
&\mathit{\Gamma}(n,\varphi) := 5m^2rt\cdot\mathrm{e}^{-(mrt)^2} + \\
&\sqrt{\frac{\varphi^2}{\varphi^2-1}}\frac{10mr}{\|\mathbf{d}\|}\mathrm{e}^{-\frac{\varphi^2 m}{\varphi^2-1}} + 2\left(\frac{2\mathrm{e}^{-1/6}}{\sqrt{3}}\right)^m + 18\varphi^2\sqrt{\frac{m}{\varphi^2-1}}\cdot\mathrm{e}^{-\frac{\varphi^4 m}{400(\varphi^2-1)}}
\end{aligned}
$$

References

[1] Aragon, N., Barreto, P.S.L.M., Bettaieb, S., Bidoux, L., Blazy, O., Deneuville, J.C., Gaborit, P., Ghosh, S., Gueron, S., Güneysu, T., Melchor, C.A., Misoczk, R., Persichett, E., Richter-Brockmann, J., Sendrie, N., Tillich, J.P., Vasseur, V., Zémor, G.: Bike: Bit flipping key encapsulation (2022), https://bikesuite.org/files/v5.0/BIKE_Spec.2022.10.10.1.pdf

[2] Aranha, D.F., Baum, C., Gjøsteen, K., Silde, T.: Verifiable mix-nets and distributed decryption for voting from lattice-based assumptions. In: Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security. p. 1467–1481. CCS '23, Association for Computing Machinery, New York, NY, USA (2023). https://doi.org/10.1145/3576915.3616683, https://doi.org/10.1145/3576915.3616683

[3] Aranha, D.F., Baum, C., Gjøsteen, K., Silde, T., Tunge, T.: Lattice-based proof of shuffle and applications to electronic voting. In: Paterson, K.G. (ed.) Topics in Cryptology – CT-RSA 2021. pp. 227–251. Springer International Publishing, Cham (2021)

[4] Asharov, G., Jain, A., López-Alt, A., Tromer, E., Vaikuntanathan, V., Wichs, D.: Multiparty computation with low communication, computation and interaction via threshold fhe. In: Pointcheval, D., Johansson, T. (eds.) Advances in Cryptology – EUROCRYPT 2012. pp. 483–501. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)

[5] Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. Mathematische Annalen **296**, 625–635 (1993). https://doi.org/10.1007/BF01445125, https://doi.org/10.1007/BF01445125

[6] Bendlin, R., Damgård, I.: Threshold decryption and zero-knowledge proofs for lattice-based cryptosystems. In: Micciancio, D. (ed.) Theory of Cryptography. pp. 201–218. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)

[7] Bindel, N., Hamburg, M., Hövelmanns, K., Hülsing, A., Persichetti, E.: Tighter proofs of CCA security in the quantum random oracle model. In: Hofheinz, D., Rosen, A. (eds.) Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part II. Lecture Notes in Computer Science, vol. 11892, pp. 61–90. Springer (2019). https://doi.org/10.1007/978-3-030-36033-7_3, https://doi.org/10.1007/978-3-030-36033-7_3

[8] Boneh, D., Gennaro, R., Goldfeder, S., Jain, A., Kim, S., Rasmussen, P.M.R., Sahai, A.: Threshold cryptosystems from threshold fully homomorphic encryption. In: Shacham, H., Boldyreva, A. (eds.) Advances in Cryptology – CRYPTO 2018. pp. 565–596. Springer International Publishing, Cham (2018)

[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)

[10] Boudgoust, K., Scholl, P.: Simple threshold (fully homomorphic) encryption from lwe with polynomial modulus. Cryptology ePrint Archive, Paper 2023/016 (2023), https://eprint.iacr.org/2023/016, https://eprint.iacr.org/2023/016

[11] Don, J., Fehr, S., Majenz, C., Schaffner, C.: Online-extractability in the quantum random-oracle model. In: Dunkelman, O., Dziembowski, S. (eds.) Advances in Cryptology – EUROCRYPT 2022. pp. 677–706. Springer International Publishing, Cham (2022)

[12] Genise, N., Micciancio, D., Peikert, C., Walter, M.: Improved discrete gaussian and subgaussian analysis for lattice cryptography. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12110, pp. 623–651. Springer (2020). https://doi.org/10.1007/978-3-030-45374-9_21, https://doi.org/10.1007/978-3-030-45374-9_21

[13] 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

[14] Hövelmanns, K., Hülsing, A., Majenz, C.: Failing gracefully: Decryption failures and the fujisaki-okamoto transform **13794**, 414–443 (2022). https://doi.org/10.1007/978-3-031-22972-5_15, https://doi.org/10.1007/978-3-031-22972-5_15

[15] Hövelmanns, K., Kiltz, E., Schäge, S., Unruh, D.: Generic authenticated key exchange in the quantum random oracle model. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part II. Lecture Notes in Computer Science, vol. 12111, pp. 389–422. Springer (2020). https://doi.org/10.1007/978-3-030-45388-6_14, https://doi.org/10.1007/978-3-030-45388-6_14

[16] Hövelmanns, K., Majenz, C.: A note on failing gracefully: Completing the picture for explicitly rejecting fujisaki-okamoto transforms using worst-case correctness. Cryptology ePrint Archive, Paper 2023/1811 (2023), https://eprint.iacr.org/2023/1811, https://eprint.iacr.org/2023/1811

[17] 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

[18] Kuchta, V., Sakzad, A., Stehlé, D., Steinfeld, R., Sun, S.: Measure-rewind-measure: Tighter quantum random oracle model proofs for one-way to

hiding and CCA security. IACR Cryptol. ePrint Arch. p. 454 (2021), https://eprint.iacr.org/2021/454

[19] Melchor, C.A., Aragon, N., Bettaieb, S., Bidoux, L., Blazy, O., Bos, J., Richter-Brockmann, J., Dion, A., Gaborit, P., Lacan, J., Persichett, E., Robert, J.M., Véron, P., Zémor, G.: Hamming quasi-cyclic (hqc) (2023), https://pqc-hqc.org/doc/hqc-specification_2023-04-30.pdf

[20] Micciancio, D., Regev, O.: Worst-case to average-case reductions based on gaussian measures. SIAM J. Comput. **37**(1), 267–302 (2007). https://doi.org/10.1137/S0097539705447360, https://doi.org/10.1137/S0097539705447360

[21] NIST: National institute for standards and technology. postquantum crypto project (2017), http://csrc.nist.gov/groups/ST/post-quantum-crypto/

[22] Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: Wagner, D.A. (ed.) Advances in Cryptology - CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings. Lecture Notes in Computer Science, vol. 5157, pp. 554–571. Springer (2008). https://doi.org/10.1007/978-3-540-85174-5_31, https://doi.org/10.1007/978-3-540-85174-5_31

[23] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM **56**(6), 34:1–34:40 (2009)

[24] Saito, T., Xagawa, K., Yamakawa, T.: Tightly-secure key-encapsulation mechanism in the quantum random oracle model. In: Nielsen, J.B., Rijmen, V. (eds.) Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III. Lecture Notes in Computer Science, vol. 10822, pp. 520–551. Springer (2018). https://doi.org/10.1007/978-3-319-78372-7_17, https://doi.org/10.1007/978-3-319-78372-7_17

[25] of Standards, N.I., Technology: Module-lattice-based key-encapsulation mechanism standard. In: FIPS 203 (2023), https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.203.ipd.pdf

[26] Targhi, E.E., Unruh, D.: Post-quantum security of the fujisaki-okamoto and oaep transforms. In: Hirt, M., Smith, A. (eds.) Theory of Cryptography. pp. 192–216. Springer Berlin Heidelberg, Berlin, Heidelberg (2016)