Module Learning with Errors with Truncated Matrices

Katharina Boudgoust[1] and Hannah Keller[2]

[1] CNRS, Univ Montpellier, LIRMM, France katharina.boudgoust@lirmm.fr
[2] Aarhus University, Denmark hkeller@cs.au.dk

Abstract. The Module Learning with Errors (MLWE) problem is one of the most commonly used hardness assumption in lattice-based cryptography. In its standard version, a matrix $\mathbf{A}$ is sampled uniformly at random over a quotient ring R_q, as well as noisy linear equations in the form of $\mathbf{As} + \mathbf{e} \bmod q$, where $\mathbf{s}$ is the secret, sampled uniformly at random over R_q, and $\mathbf{e}$ is the error, coming from a Gaussian distribution. Many previous works have focused on variants of MLWE, where the secret and/or the error are sampled from different distributions. Only few works have focused on different distributions for the matrix $\mathbf{A}$. One variant proposed in the literature is to consider matrix distributions, where the low-order bits of a uniform $\mathbf{A}$ are deleted. This seems a natural approach in order to save in bandwidth. We call it *truncated* MLWE.
In this work, we show that the hardness of standard MLWE implies the hardness of truncated MLWE, both for search and decision versions. Prior works only covered the search variant and relied on the (module) NTRU assumption, limitations which we are able to overcome. Overall, we provide two approaches, offering different advantages. The first uses a general Rényi divergence argument, applicable to a wide range of secret/error distributions, but which only works for the search variants of (truncated) MLWE. The second applies to the decision versions, by going through an intermediate variant of MLWE, where additional *hints* on the secret are given to the adversary. However, the reduction makes use of discrete Gaussian distributions.

Keywords: Lattices, Module Learning with Errors, Truncation

1 Introduction

The Module Learning with Errors (MLWE) problem [LS15] is among the most commonly used hardness assumptions in lattice-based cryptography. Besides its strong connection to well-studied, worst-case, structured lattice problems, it also comes with an easy-to-work-with shape in the language of linear algebra. It has shown to be very versatile in its possible applications in cryptography. Informally, MLWE can be seen as noisy linear equations over the quotient $R_q := R/qR$ for some ring R and positive integer q. More formally, a sample of MLWE can be described as follows. Sample a *matrix* $\mathbf{A}$ from a distribution $\mathsf{D}_{\mathsf{mat}}$ over R_q, a

secret vector $\mathbf{s}$ from a distribution $\mathsf{D}_{\mathsf{sec}}$ over R_q and an *error vector* $\mathbf{e}$ from a distribution $\mathsf{D}_{\mathsf{err}}$ over R. Then, compute $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$ and output $(\mathbf{A}, \mathbf{b})$. The search variant of MLWE asks to find the secret $\mathbf{s}$, whereas the decision variant asks to distinguish an MLWE sample from an instance of the uniform distribution of matrices and vectors over R_q. Originally, the problem was studied over the special ring $R = \mathbb{Z}$, and termed Learning with Errors (LWE) [Reg05,Reg09]. Later, the problem was generalized to the ring of integers of number fields of higher degrees [LS15]. The original formulation of MLWE, which is connected by a worst-case to average-case reduction to well-studied module lattice problems, sets $\mathsf{D}_{\mathsf{mat}}$ and $\mathsf{D}_{\mathsf{sec}}$ as the uniform distributions over R_q and $\mathsf{D}_{\mathsf{err}}$ as a rounded or discrete Gaussian distribution [Reg05,LS15]. The reduction first made use of quantum algorithms, but was later made classical [Pei09,BLP+13,BJRW20].

Since then, different lines of work studied the hardness of MLWE for different distributions. Regarding variants for the secret distribution, an early result showed that, with only a small loss in the row dimension of $\mathbf{A}$, the secret distribution $\mathsf{D}_{\mathsf{sec}}$ can be set the same as the error distribution $\mathsf{D}_{\mathsf{err}}$ [ACPS09]. This variant is commonly referred to as MLWE in its *Hermite normal form*. Moreover, the hardness of MLWE where the secret is sampled uniformly over a small subset of R_q was established for the degree-1 case in [GKPV10,BLP+13,Mic18] (focusing on the special subset $\{0, 1\}$). It was then generalized to rings of larger degrees [BJRW20,BJRW23] and to *any* secret distribution with enough min-entropy [BD20,BJRW22,LWZW24]. Regarding variants for the error distribution, different results have shown the hardness of MLWE if the error is sampled uniformly over a small subset of $\mathbb{Z}_q$ [DM13,MP13,BCD+16,BLR+18,STA20] and for higher-degree rings R_q [BJRW23].

So far, only few works have studied the hardness of MLWE when the matrix $\mathbf{A}$ does not follow the uniform distribution over R_q.[3] By a rather simple reduction, one can reduce standard MLWE with a uniform matrix to a variant where $\mathbf{A}$ is composed of polynomials which only have binary coefficients.[4] The idea is to compute the bit-decomposition of every coefficient of each polynomial entry in the uniform $\mathbf{A} = \mathsf{bin}(\mathbf{A}) \cdot \mathbf{G}$, where $\mathbf{G}$ is the so-called gadget matrix. A given MLWE instance $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ then automatically defines an instance of MLWE with a binary matrix $(\mathsf{bin}(\mathbf{A}), \mathsf{bin}(\mathbf{A})\mathbf{s}' + \mathbf{e})$, where $\mathbf{s}' = \mathbf{G}\mathbf{s}$. If $\mathbf{A}$ was originally an $m \times n$ matrix, $\mathsf{bin}(\mathbf{A})$ is now an $m \times (n \cdot \lceil \log q \rceil)$ matrix. This approach has been formalized and generalized for plain LWE in [BLMR13] to what they call *coset sampleable* distributions. Besides binary distributions, they cover discrete Gaussians and uniform distribution over linear subspaces. Other works have used matrix distributions $\mathsf{D}_{\mathsf{mat}}$ that are computationally [GKPV10,BD20], statistically [Reg05,GPV08], or Rényi [BLR+18] close to the uniform distribution. In a recent work [JLS24], the study of plain LWE with a sparse matrix was initiated, yielding improved computation and storage efficiency.

[3] This is called Non-uniform Learning with Errors in [BLMR13].

[4] Throughout this work, we use the so-called coefficient embedding to identify elements in R_q with polynomials having coefficients in $\mathbb{Z}_q$.

Truncated MLWE. In this work, we study the hardness of MLWE for a different matrix distribution which has been considered in a recent result [JZW+23].[5] The formulation of the problem is rather simple. To sample the matrix, for some small constant c, one samples some matrix $\mathbf{U}$ uniformly at random over R_q, then deletes the c lowest-order bits of every coefficient of each entry in $\mathbf{U}$. We say that the matrix is *truncated* and write $\mathbf{A} = \mathsf{Trunc}(\mathbf{U}, c)$. As before, a sample is given by $(\mathbf{A}, \mathbf{As} + \mathbf{e})$ for some secret $\mathbf{s}$ and error $\mathbf{e}$. Subsequently, we call the variant the *truncated* MLWE problem. Intuitively, the motivation of this variant is to save in bandwidth. Whenever we have to send an MLWE instance $(\mathbf{A}, \mathbf{b})$ (in form of a public key or an encrypted message, for instance), the size of the message to be sent is smaller if we delete the low-order bits of every entry of the matrix. In [JZW+23], a reduction from the module variant of the NTRU problem to the search variant of truncated MLWE (with entropic secret) was proven. As the NTRU assumption [HPS98] and its module version [CPS+20] are seen as less standard than MLWE and the search problem is not enough for many security notions, like standard IND-CPA security of encryption schemes, we would ideally like to show that the hardness of decision truncated MLWE can be reduced from the hardness of standard MLWE. This leaves the following open problem stated by [JZW+23], motivating our work:

> *Does the hardness of standard* MLWE *imply the hardness of search and decision truncated* MLWE*?*

Our Contributions. We answer this research question positively. We show two approaches for how the hardness of standard MLWE implies the hardness of truncated MLWE, both for the search and decision variants. Each of the approaches comes with different advantages. We provide a detailed comparison between our two proofs and the results of [JZW+23] in Section 6.

First Approach. In Section 4, we reduce the hardness of truncated MLWE from standard MLWE using the Rényi divergence as a measure of distance. The Rényi divergence has been used for tight reductions in lattice-based cryptography since [BLL+15,BLR+18]. The high level idea of Theorem 2 is to view a truncated matrix $\mathbf{A} = \mathsf{Trunc}(\mathbf{U}, c)$ as the difference of the original uniform matrix $\mathbf{U}$ and the deleted low-order bits $\mathbf{N_U}$, i.e., $\mathbf{A} = \mathbf{U} - \mathbf{N_U}$. Then a sample $(\mathsf{Trunc}(\mathbf{U}, c), \mathsf{Trunc}(\mathbf{U}, c)\mathbf{s} + \mathbf{e})$ can be viewed as an instance of standard MLWE given by $(\mathbf{U}, \mathbf{Us} + \mathbf{e}')$, where $\mathbf{e}' = -\mathbf{N_U s} + \mathbf{e}$. Note that $\mathbf{e}'$ currently depends on the secret $\mathbf{s}$ and might thus leak sensitive information about it. By a standard Rényi argument, one can make the distribution of $\mathbf{e}'$ independent of $\mathbf{N_U s}$, as long as the error distribution is sufficiently large. The resulting loss in advantage depends on the ring degree, the size of elements coming from the secret distribution $\mathsf{D}_{\mathsf{sec}}$, the number of deleted bits c, the dimensions of the matrix, as well as the error distribution $\mathsf{D}_{\mathsf{err}}$. The result generally applies to any secret

[5] In [JZW+23], a more general notion of MLWE with *semiuniform matrices* is introduced. As we are not aware of any concrete applications of their more general notion, we decided to keep the presentation of the problem as simple as possible in our work.

and noise distributions for which MLWE is believed to be hard, as long as we can compute the relevant Rényi divergences, but is restricted to the corresponding search variant of the problems. Recent results have for instance put forward the use of Rényi divergence arguments in combination with sum of bounded uniform distributions [dPKPR24]. In contrast to discrete Gaussian distributions, they are easier to implement and to protect against side-channel attacks. As of today, the only way to use Rényi divergence arguments for decision variants, is to make use of the so-called *public sampleability framework* of [BLR+18]. However, as we argue in Section 4.1, this framework only leads to a vacuous reduction in our context, as the Rényi divergence between truncated and non-truncated matrices is exponentially large in their dimensions.

Second Approach. To circumvent this issue, we propose an alternative approach in Section 5, covering both the search and decision versions. The main idea is to interpret the information $\mathbf{N_U s}$ leaked about the secret $\mathbf{s}$ as approximate *hints*. The presence of hints is defining another (already studied) variant of MLWE, whose hardness can be derived from standard MLWE, both for the decision and search variants [MKMS22,KLSS23]. Informally, the type of hints we are considering is $\mathbf{Hs} + \mathbf{f}$, for some hint matrix $\mathbf{H}$ known to the adversary and some noise term $\mathbf{f}$ unknown to the adversary. In the truncated context, we can simply set $\mathbf{H}$ to store the low-order bits of the matrix, i.e., $\mathbf{H} = \mathbf{N_U}$. We recall the formal definition and (a generalized) hardness results of MLWE with hints in Section 3.2 and then show a reduction from MLWE with hints to truncated MLWE in Theorem 3. The advantage of now applying this result to the decision variant comes with the drawback that the hardness results of MLWE with hints only apply to a limited set of secret and noise distributions. More precisely, the existing reductions make use of decomposition theorems for discrete Gaussian distributions.

Choice of Rings. All of our results are proven for the class of power-of-two cyclotomic rings. This restriction is mainly due to tighter reductions, as we have a good control over the norm growth after multiplying two elements (interpreted as polynomials) in such rings. It is possible to generalize everything to other fields and rings, incurring some additional reduction losses due to the so-called expansion factor [LM06,RSW18]. As power-of-two cyclotomic rings are the most popular choice, both in theory and in practice, we opted for directly showing the tighter results.

Trivial Setup. We would like to mention that there is a setup of truncated MLWE which makes it trivially easy to solve. Let c be the number of bits we are truncating away from the matrix, i.e, $\mathbf{A} = \mathsf{Trunc}(\mathbf{U}, c) = \mathbf{U} - \mathbf{N_U}$. If 2^c is a factor of the modulus q, we know that $(\mathbf{As} \bmod q) \bmod 2^c = \mathbf{0}$ for every MLWE secret $\mathbf{s}$. If additionally the noise $\mathbf{e}$ has infinity norm less than 2^c, it would be easy to solve the truncated MLWE instance. On input $(\mathbf{A}, \mathbf{b})$ with $\mathbf{b} = \mathbf{As} + \mathbf{e} \bmod q$, we can simply compute $\mathbf{b} \bmod 2^c$ to recover $\mathbf{e}$. We highlight that our reductions do not allow for this trivial setup. The reduction of Theorem 2 in Section 4 requires the resulting error distribution of truncated MLWE to be significantly larger than the

shift (that is, $\mathbf{N_U s}$) it is trying to hide. This shift is (among other parameters) determined by 2^c, so the error distribution cannot have infinity norm below 2^c. Similarly, the reduction of Theorem 1 in Section 5 requires the resulting error distribution to be significantly (among other parameters) larger than the infinity bound on the hint matrix $\mathbf{N_U}$, which is bounded by 2^c.

Future Directions. Our approach can be seen as an additive decomposition of a uniform matrix $\mathbf{U}$ into a matrix containing the low-order bits $\mathbf{N_U}$ and a matrix containing the high-order bits $\mathsf{Trunc}(\mathbf{U}, c)$. The coset sampleable approach in [BLMR13], on the other hand, uses a multiplicative decomposition of a uniform matrix $\mathbf{U}$ into a (fixed) matrix $\mathbf{M}$ and a matrix $\mathbf{V}$ coming from the aimed distribution (e.g., binary or discrete Gaussian). One future direction could be to capture both the additive and multiplicative decompositions in one result.

2 Preliminaries

2.1 Notations

For any positive integer q, we denote by $\mathbb{Z}_q := \mathbb{Z}/q\mathbb{Z}$ the quotient integer ring. Elements in $\mathbb{Z}$ can be reduced mod q and possess a unique representative in the set $\{0, \ldots, q-1\}$. Column vectors are written in bold lowercase letters $\mathbf{b}$ and matrices in bold uppercase letters $\mathbf{A}$. The transpose operator over vectors and matrices is denoted by $\mathbf{b}^T$ and $\mathbf{A}^T$. The determinant of a matrix $\mathbf{A}$ is denoted by $\det(\mathbf{A})$. For any vector $\mathbf{b}$, we denote by $\|\mathbf{b}\|$ its ℓ_2-norm and by $\|\mathbf{b}\|_\infty$ its infinity norm. For any matrix $\mathbf{A}$, we denote by $\|\mathbf{A}\|_\infty$ the maximum of the infinity norms of its column vectors. For any real number $r \in \mathbb{R}$, the operation $\lfloor r \rceil$ denotes rounding it to the nearest integer (with 0.5 being rounded up). The operation $\lceil r \rceil$ denotes rounding it up to the next integer. We can component-wise extend rounding to vectors and matrices. All logarithms are base 2. By $\mathsf{negl}(\lambda)$ we denote a negligible function in λ, thus it decreases faster towards 0 than the inverse of any polynomial function. The abbreviation PPT stands for probabilistic polynomial-time.

We define a truncation function Trunc which takes as input an element x in $\mathbb{Z}_q$ and a positive integer c, computes and outputs an element in $\mathbb{Z}_q$:

$$\mathsf{Trunc}(x, c) = x - (x \bmod 2^c). \tag{1}$$

Informally, during truncation the c lowest bits of $x \in R_q$ are set to 0. We can naturally extend the truncation function to vectors and matrices over $\mathbb{Z}$ by applying them coefficient-wise and entry-wise, respectively.

Let n be a positive integer. An $n \times n$ symmetric real matrix $\mathbf{M}$ is said to be *positive semidefinite* if $\mathbf{x}^T\mathbf{M}\mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$. Moreover, an $n \times n$ matrix $\mathbf{M} = (m_{ij})_{i,j \in \{1,\ldots,n\}}$ is called *diagonally dominant* if $|m_{ii}| \geq \sum_{j \neq i} |m_{ij}|$ for all $i \in \{1, \ldots, n\}$. A symmetric diagonally dominant matrix with real non-negative diagonal entries is positive semidefinite.

2.2 Number Theory

A number field $K = \mathbb{Q}(\zeta)$ of degree d is a finite field extension of the rationals $\mathbb{Q}$ obtained by adjoining an algebraic number ζ. We denote its ring of integers by R. We call K a ν-th cyclotomic number field if ζ is a ν-th primitive root of unity. Its degree is given by $d = \varphi(\nu)$, where φ is Euler's totient function. We say R is a power-of-two cyclotomic, if it is the ring of integers of the ν-th cyclotomic field, where ν can be written as 2^{k+1} for some positive integer k. In that case, $d = 2^k$.

We can identify $K = \mathbb{Q}[X]/\langle \Phi(X) \rangle$, where $\Phi(X)$ is the minimal polynomial of ζ. Every element $x \in K$ can then by written with respect to the basis $\left\{1, \zeta, \ldots, \zeta^{d-1}\right\}$, thus $x = \sum_{i=0}^{d-1} x_i \zeta^i$ with $x_i \in \mathbb{Q}$. The isomorphism $\tau \colon K \to \mathbb{Q}^d$ which maps x to its coefficient vector $\tau(x) = (x_0, \ldots, x_{d-1})^T$ is called the coefficient embedding. By restricting τ to R, we obtain an isomorphism between R and $\mathbb{Z}^d$. By associating the norm of an element x in R with the norm of its corresponding $\tau(x) \in \mathbb{Z}^d$, it is possible to equip R with a geometry. With this geometry at hand, we can define norms of vectors and matrices over R, as well as round and truncate elements in R coefficient-wise.

Every product of two ring elements $x \cdot y = z \in R$ can be represented as a matrix vector product over $\mathbb{Z}$, such that $\mathsf{Rot}(x) \cdot \tau(y) = \tau(z) \in \mathbb{Z}^d$. We call $\mathsf{Rot}(x)$ the rotation matrix associated to x in the coefficient embedding. The exact shape of $\mathsf{Rot}(x)$ depends on the number field (and associated ring of integers) we are considering. Throughout the paper, we make use of the fact that for power-of-two cyclotomics, $\mathsf{Rot}(x)$ is nega-cyclic and that $\|\mathsf{Rot}(x)\|_\infty = \|x\|_\infty$. One could move to different fields and rings, at the expenses of the norm of the rotation matrix being larger than the norm of the underlying ring element. The multiplicative factor is sometimes also called expansion factor of K [LM06,RSW18].

2.3 Lattices

Let d be a positive integer. A (full-rank) Euclidean lattice Λ is a discrete subgroup of $\mathbb{R}^d$ and can be represented by some basis vectors $\{\mathbf{b}_1, \ldots, \mathbf{b}_d\} \in \mathbb{R}^d$ s.t. $\Lambda = \left\{\sum_{i=1}^{d} z_i \mathbf{b}_i \mid z_i \in \mathbb{Z}\right\}$. Let $\mathbf{B} = (\mathbf{b}_i)_{i \in \{1,\ldots,d\}}$ be the matrix composed of the basis column vectors. The determinant of a lattice is defined as $\det(\Lambda) = \det(\mathbf{B})$. We further define the span of a lattice as $\mathsf{Span}(\Lambda) = \left\{\sum_{i=1}^{d} r_i \mathbf{b}_i \mid r_i \in \mathbb{R}\right\}$ and its dual by $\Lambda^* = \left\{\mathbf{x} \in \mathsf{Span}(\Lambda) \mid \mathbf{x}^T \mathbf{y} \in \mathbb{Z}\ \forall \mathbf{y} \in \Lambda\right\}$. Every ring of integers R defines a lattice $\Lambda = \{\tau(x) \mid x \in R\}$, using the coefficient embedding.

2.4 Probability Measures

For a finite set S, we denote by $x \leftarrow S$ the process of sampling x uniformly at random over S. For d, and $k_1 \leq k_2$ positive integers, let $\mathbf{U}_{k_1,k_2}$ denote the uniform distribution over $\{k_1, \ldots, k_2 - 1\}^d$, i.e., $\Pr_{\mathbf{x} \leftarrow \mathbf{U}_{k_1,k_2}}[\mathbf{x} = \mathbf{y}] = (1/(k_2 - k_1))^d$ for every $\mathbf{y} \in \{k_1, \ldots, k_2 - 1\}^d$. By using the coefficient embedding $\tau \colon R \to \mathbb{Z}^d$, this can be seen as a distribution over R.

Definition 1. *Let $B, \delta > 0$ and R be a degree d ring of integers. A distribution $\mathcal{D}$ over R is (B, δ)-bounded if*

$$\Pr_{x \leftarrow \mathcal{D}}[\|\tau(x)\|_\infty > B] \leq \delta,$$

where τ is the coefficient embedding of R into $\mathbb{Z}^d$.

Continuous, Rounded and Discrete Gaussian Distributions.

Definition 2 (One-dimensional Gaussian Distribution). *Probability distribution D_{μ,σ^2} with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 \in \mathbb{R}$ samples value $x \in \mathbb{R}$ with probability distribution function*

$$D_{\mu,\sigma^2}(x) := \frac{1}{\sigma\sqrt{2\pi}} \exp(-(x-\mu)^2/(2\sigma^2)).$$

Definition 3 (Multivariate Gaussian Distribution). *Probability distribution $D_{\boldsymbol{\mu},\boldsymbol{\Sigma}}$ with mean $\boldsymbol{\mu} \in \mathbb{R}^d$ and covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{d\times d}$ samples vector $\mathbf{x} \in \mathbb{R}^d$ with probability distribution function*

$$D_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\mathbf{x}) := \frac{1}{\sqrt{(2\pi)^d \det(\boldsymbol{\Sigma})}} \exp(-(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}-\boldsymbol{\mu})^T/2).$$

If $\boldsymbol{\Sigma} = \sigma^2 \cdot \mathbf{I}_d$, we call the distribution spherical and simply write $D_{\boldsymbol{\mu},\sigma^2}$. For $\boldsymbol{\mu} = \mathbf{0}$, we might omit it from the notation and simply write $D_{\boldsymbol{\Sigma}}$.

We further define the *rounded* Gaussian distribution $\lfloor D_{\boldsymbol{\mu},\boldsymbol{\Sigma}} \rceil$ over $\mathbb{Z}^d$, where the instance sampled from the continuous Gaussian distribution over $\mathbb{R}^d$ is rounded to the nearest integer.

Definition 4 (Discrete Gaussian Distribution for Lattices). *Probability distribution $D_{\Lambda,\boldsymbol{\mu},\boldsymbol{\Sigma}}$ over a lattice $\Lambda \subseteq \mathbb{R}^d$ with mean $\boldsymbol{\mu} \in \mathbb{R}^d$ and covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{d\times d}$ samples vector $\mathbf{x} \in \Lambda \subseteq \mathbb{R}^d$ in lattice Λ with probability distribution function*

$$D_{\Lambda,\boldsymbol{\mu},\boldsymbol{\Sigma}}(\mathbf{x}) := D_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\mathbf{x}) / \sum_{\mathbf{y}\in\Lambda} D_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\mathbf{y}).$$

As for the continuous case, if $\boldsymbol{\Sigma} = \sigma^2 \cdot \mathbf{I}_d$, we simply write $D_{\Lambda,\boldsymbol{\mu},\sigma^2}$. For $\boldsymbol{\mu} = \mathbf{0}$, we might omit it from the notation and simply write $D_{\Lambda,\boldsymbol{\Sigma}}$.

The smoothing parameter of a lattice Λ, denoted by $\eta_\epsilon(\Lambda)$ for some $\epsilon > 0$ and introduced by [MR04], is the smallest $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \epsilon$, where $\rho_\sigma(\mathbf{x}) := \exp(-\pi\|\mathbf{x}\|^2/\sigma^2)$. When ϵ is omitted, it is some unspecified negligible function $\epsilon = \mathsf{negl}(\lambda)$ in the lattice dimension or the security parameter. By specializing [MR04, Lem. 3.2] to the integer lattice $\Lambda = \mathbb{Z}^d$ (which is self-dual), we know that for $\epsilon = 2^{-d}$ it holds $\eta_\epsilon(\Lambda) \leq \sqrt{d}$.

We use the coefficient embedding τ to sample discrete Gaussian distributions over R of degree d. We denote by $\mathbf{s} \leftarrow D_{R^m,\boldsymbol{\mu},\boldsymbol{\Sigma}}$ the process of sampling $\mathbf{s}' \leftarrow D_{\mathbb{Z}^{dm},\boldsymbol{\mu},\boldsymbol{\Sigma}}$ and setting $\mathbf{s} := \tau^{-1}(\mathbf{s}')$.

Lemma 1 (Adapted from [Lyu12, Lem. 4.4]). *Let t, σ be positive reals and R be a degree-d ring of integers. Then D_{R,σ^2} is $(t, 2d\exp(-t^2/2\sigma^2))$-bounded.*

Lemma 2 (Adapted from [MR04, Lem. 4.4]). *Let Λ be an n-dimensional lattice and $\epsilon \in (0,1)$. Then for any $\boldsymbol{c} \in \mathbb{R}^n$ and $\sigma \geq \eta_\epsilon(\Lambda)$ we have that $\rho_{\mathbf{c},\sigma^2}(\Lambda) := \sum_{\mathbf{x}\in\Lambda} \exp(-\pi\|\mathbf{x}-\boldsymbol{c}\|^2/\sigma^2)$ is in the range $[1-\epsilon, 1+\epsilon]\cdot\det(\Lambda)^{-1}$.*

The smoothing parameter is relevant when decomposing discrete Gaussians.

Lemma 3 (Decomposition).

1. *Let $\sigma, \delta \in \mathbb{R}$ be two variances and $\Lambda \subset \mathbb{R}^d$ be a lattice. Let $\mathbf{x}_1 \leftarrow D_{\Lambda,\sigma^2}$ and $\mathbf{x}_2 \leftarrow D_{\Lambda,\delta^2}$ with $\sigma, \delta \geq \sqrt{2}\cdot\eta(\Lambda)$. Then, $\mathbf{x} := \mathbf{x}_1 + \mathbf{x}_2$ is statistically close to a zero-centered discrete Gaussian distribution over Λ with covariance $\gamma^2 = \sigma^2 + \delta^2$.*
2. *Let $\Lambda \subset \mathbb{Z}^m$ be a sub-lattice of rank n with basis $\mathbf{L} \in \mathbb{Z}^{m\times n}$. Further, let σ be a positive real, defining $\boldsymbol{\Sigma} = \sigma^2\mathbf{L}\mathbf{L}^T \in \mathbb{Z}^{m\times m}$, and $\boldsymbol{\Sigma}' \in \mathbb{Z}^{m\times m}$ be a positive semidefinite matrix. Moreover, we assume that the eigenvalues of the matrix $\Gamma = \sigma\sqrt{\mathbf{I}_m - \sigma^2\mathbf{L}(\boldsymbol{\Sigma}+\boldsymbol{\Sigma}')^{-1}\mathbf{L}^T}$ are greater than or equal to the smoothing parameter $\eta(\mathbb{Z}^m)$. Let $\mathbf{x}_1 \leftarrow D_{\mathbb{Z}^m,\boldsymbol{\Sigma}'}$ and $\mathbf{x}_2 \leftarrow D_{\Lambda,\boldsymbol{\Sigma}}$. Then, $\mathbf{x} := \mathbf{x}_1 + \mathbf{x}_2$ is statistically close to a zero-centered discrete Gaussian distribution over $\mathbb{Z}^m$ with covariance matrix $\boldsymbol{\Sigma}+\boldsymbol{\Sigma}'$.*

Proof. Item 1. This is a special case of [MP13, Thm. 3.3] with $m = 2$ and $\mathbf{z}$ the all-1 vector. Item 2. Proven in [MKMS22, Lem. 1]. □

Measurement of Distribution Closeness. In the following, we recall the definition of the Rényi divergence of some order α. Even though it is technically possible to allow the order to take a real value, we limit ourselves to α being a positive integer throughout the work.

Definition 5 (Rényi Divergence). *For any two probability distributions P and Q defined over $\mathbb{R}$, the Rényi divergence (RD) of order $\alpha > 1$ is defined as*

$$\mathsf{RD}_\alpha(P\|Q) = \mathbb{E}_{x\sim Q}\left(\frac{P(x)}{Q(x)}\right)^{\frac{\alpha}{\alpha-1}}.$$

For discrete distributions with $\mathsf{Supp}(P) \subseteq \mathsf{Supp}(Q)$, this is:

$$\mathsf{RD}_\alpha(P\|Q) = \left(\sum_{x\in\mathsf{Supp}(P)} \frac{P(x)^\alpha}{Q(x)^{\alpha-1}}\right)^{\frac{1}{\alpha-1}}.$$

Lemma 4 (Multiplicativity [LSS14, Lemma 4.1]). *Let $\alpha \in (1,\infty)$. Let P and Q denote distributions of a pair of random variables (Y_1, Y_2). Also, for $i \in \{1,2\}$ let P_i and Q_i be the marginal distribution of Y_i under P and Q, respectively. Then if Y_1 and Y_2 are independent:*

$$\mathsf{RD}_\alpha(P\|Q) = \mathsf{RD}_\alpha(P_1\|Q_1)\cdot\mathsf{RD}_\alpha(P_2\|Q_2).$$

Lemma 5 (Probability Preservation [LSS14, Lemma 4.1]). *Let $\alpha \in (1,\infty)$ and $E \subseteq \mathsf{Supp}(Q)$ be an arbitrary event. Then:*

$$Q(E) \geq P(E)^{\frac{\alpha}{\alpha-1}} / \mathsf{RD}_\alpha(P\|Q).$$

Lemma 6 (Data Processing Inequality [vEH14, Theorem 9]). *Let $\alpha \in (1,\infty)$. For any function f, where P^f (respectively Q^f) denotes the distribution of $f(y)$ induced by sampling $y \leftarrow P$ (respectively $y \leftarrow Q$):*

$$\mathsf{RD}_\alpha(P^f\|Q^f) \leq \mathsf{RD}_\alpha(P\|Q).$$

Lemma 7 ([Mir17, Prop. 7]). *For mean $\mu \in \mathbb{R}$, variance $\sigma^2 \in \mathbb{R}$ and order $\alpha > 1$ it holds*

$$\mathsf{RD}_\alpha(D_{\sigma^2}\|D_{\mu,\sigma^2}) = \mathsf{RD}_\alpha(D_{\mu,\sigma^2}\|D_{\sigma^2}) = \exp(\alpha\mu^2/(2\sigma^2)).$$

The following lemma generalizes a result on the Rényi divergence of discrete Gaussians from [LSS14] to arbitrary orders. We specialize it to spherical discrete Gaussians for simplicity of presentation.

Lemma 8 (Adapted from [LSS14, Lem. 4.2]). *Let α be a positive integer, $\boldsymbol{\mu_1}, \boldsymbol{\mu_2} \in \mathbb{R}^d$ and σ be a positive real. Further, let $\Lambda \subset \mathbb{Z}^d$ be a lattice. If $\boldsymbol{\mu_1}, \boldsymbol{\mu_2} \in \Lambda$, let $\epsilon = 0$. Otherwise, fix $\epsilon \in (0,1)$ and assume $\sigma \geq \eta_\epsilon(\Lambda)$. For any lattice $\Lambda \in \mathbb{R}^d$:*

$$\mathsf{RD}_\alpha := \mathsf{RD}_\alpha(D_{\Lambda,\boldsymbol{\mu_1},\sigma^2}\|D_{\Lambda,\boldsymbol{\mu_2},\sigma^2}) \leq \left(\frac{1+\epsilon}{1-\epsilon}\right)^{\alpha/(\alpha-1)} \cdot \exp(\alpha\|\boldsymbol{\mu_1}-\boldsymbol{\mu_2}\|^2/(2\sigma^2)).$$

Proof. By definition of discrete Gaussians,

$$D_{\Lambda,\boldsymbol{\mu_1},\sigma^2}(\mathbf{x}) = \frac{\exp(-\|\mathbf{x}-\boldsymbol{\mu_1}\|^2/(2\sigma^2))}{\rho_{\boldsymbol{\mu_1},\sigma^2}(\Lambda)} \quad \text{and}$$
$$D_{\Lambda,\boldsymbol{\mu_2},\sigma^2}(\mathbf{x}) = \frac{\exp(-\|\mathbf{x}-\boldsymbol{\mu_2}\|^2/(2\sigma^2))}{\rho_{\boldsymbol{\mu_2},\sigma^2}(\Lambda)},$$

where $\rho_{\boldsymbol{\mu},\sigma^2}(\Lambda) = \sum_{\mathbf{y}\in\Lambda} \exp(-\|\mathbf{y}-\boldsymbol{\mu_1}\|^2/(2\sigma^2))$ for any $\boldsymbol{\mu} \in \mathbb{R}^d$. We compute

$$\mathsf{RD}_\alpha = \left(\sum_{\mathbf{x}\in\Lambda} \frac{D_{\Lambda,\boldsymbol{\mu_1},\sigma^2}(\mathbf{x})^\alpha}{D_{\Lambda,\boldsymbol{\mu_2},\sigma^2}(\mathbf{x})^{\alpha-1}}\right)^{1/(\alpha-1)} = \frac{\rho_{\boldsymbol{\mu_2},\sigma^2}(\Lambda)}{\rho_{\boldsymbol{\mu_1},\sigma^2}(\Lambda)^{\alpha/(\alpha-1)}}$$
$$\cdot \left(\sum_{\mathbf{x}\in\Lambda} \exp(-\alpha\|\mathbf{x}-\boldsymbol{\mu_1}\|^2/(2\sigma^2) + (\alpha-1)\|\mathbf{x}-\boldsymbol{\mu_2}\|^2/(2\sigma^2))\right)^{1/(\alpha-1)}.$$

We first simplify the right term of the multiplication, then simplify the whole multiplication. Defining $\boldsymbol{c} = \alpha\boldsymbol{\mu_1} - (\alpha-1)\boldsymbol{\mu_2}$ we claim that:

Claim. $\alpha\|\mathbf{x}-\boldsymbol{\mu_1}\|^2 - (\alpha-1)\|\mathbf{x}-\boldsymbol{\mu_2}\|^2 = \|\mathbf{x}-\boldsymbol{c}\|^2 - \alpha(\alpha-1)\|\boldsymbol{\mu_1}-\boldsymbol{\mu_2}\|^2$.

Proof.

$$\begin{aligned}
&\alpha\|\mathbf{x}-\boldsymbol{\mu_1}\|^2-(\alpha-1)\|\mathbf{x}-\boldsymbol{\mu_2}\|^2\\
&=\|\mathbf{x}\|^2+(\alpha-1)^2\|\boldsymbol{\mu_2}\|^2+2(\alpha-1)\langle\mathbf{x},\boldsymbol{\mu_2}\rangle+\alpha^2\|\boldsymbol{\mu_1}\|^2-2\alpha\langle\mathbf{x},\boldsymbol{\mu_1}\rangle\\
&\quad-2\alpha\langle(\alpha-1)\boldsymbol{\mu_2},\boldsymbol{\mu_1}\rangle-(\alpha-1)^2\|\boldsymbol{\mu_2}\|^2-(\alpha-1)\|\boldsymbol{\mu_2}\|^2+\alpha\|\boldsymbol{\mu_1}\|^2\\
&\quad-\alpha^2\|\boldsymbol{\mu_1}\|^2+2\alpha\langle(\alpha-1)\boldsymbol{\mu_2},\boldsymbol{\mu_1}\rangle\\
&=\|\mathbf{x}-\boldsymbol{c}\|^2-\alpha(\alpha-1)\|\boldsymbol{\mu_1}-\boldsymbol{\mu_2}\|^2
\end{aligned}$$

■

Hence the right term of the multiplication simplifies as follows:

$$\begin{aligned}
&\left(\sum_{\mathbf{x}\in\Lambda}\exp(-\alpha\|\mathbf{x}-\boldsymbol{\mu_1}\|^2/(2\sigma^2)+(\alpha-1)\|\mathbf{x}-\boldsymbol{\mu_2}\|^2/(2\sigma^2))\right)^{1/(\alpha-1)}\\
&=\exp(\alpha\|\boldsymbol{\mu_1}-\boldsymbol{\mu_2}\|^2/(2\sigma^2))\cdot\left(\sum_{\mathbf{x}\in\Lambda}\exp(-\|\mathbf{x}-\boldsymbol{c}\|^2/(2\sigma^2))\right)^{1/(\alpha-1)}\\
&=\exp(\alpha\|\boldsymbol{\mu_1}-\boldsymbol{\mu_2}\|^2/(2\sigma^2))\cdot\rho_{\mathbf{c},\sigma^2}(\Lambda)^{1/(\alpha-1)}.
\end{aligned}$$

Notice that for $\boldsymbol{\mu_1},\boldsymbol{\mu_2}\in\Lambda$ and thus $\mathbf{c}\in\Lambda$, we have $\rho_{\boldsymbol{\mu_1},\sigma^2}(\Lambda)=\rho_{\boldsymbol{\mu_2},\sigma^2}(\Lambda)=\rho_{\mathbf{c},\sigma^2}(\Lambda)$. From this, we conclude that

$$\frac{\rho_{\boldsymbol{\mu_2},\sigma^2}(\Lambda)}{\rho_{\boldsymbol{\mu_1},\sigma^2}(\Lambda)^{\alpha/(\alpha-1)}}\cdot\rho_{\mathbf{c},\sigma^2}(\Lambda)^{1/(\alpha-1)}=1.$$

As a result, we get $\mathsf{RD}_\alpha=\exp(\alpha\|\boldsymbol{\mu_1}-\boldsymbol{\mu_2}\|^2/(2\sigma^2))$.

Otherwise, we can use the assumption that $\sigma\geq\eta_\epsilon(\Lambda)$ and apply Lemma 2, fixing $\epsilon\in(0,1)$, from which we know that for any $\boldsymbol{z}\in\mathbb{R}^n$, $\rho_{\boldsymbol{z},\sigma^2}(\Lambda)$ is in the range $[1-\epsilon,1+\epsilon]\cdot\det(\Lambda)^{-1}$. Applying this to the sums in the expression for RD_α gives the claimed interval for RD_α. □

3 Module Learning with Errors and Variants

We first introduce in Section 3.1 the definition of truncated Module Learning with Errors, from which standard LWE [Reg05] and MLWE [LS15] can be obtained as special cases. In Section 3.2, we recall the definition of MLWE with hints on the secret and prove that its hardness can be reduced from standard MLWE.

3.1 Truncated Module Learning with Errors

Truncated MLWE as we define it has a matrix $\mathbf{A}\in R_q^{m\times n}$ whose entries are truncated, where Trunc is defined in Equation 1. The only existing similar definition in the literature is given in [JZW+23], where truncation can be seen as a special case of semiuniform distributions.

Definition 6 (Truncated MLWE Problem). *Let R be a degree-d ring of integers. Let q, m, n and c be positive integers. Further let $\mathsf{D}_{\mathsf{err}}$ and $\mathsf{D}_{\mathsf{sec}}$ be distributions over R. The (non-interactive) experiments of the search and decision versions of truncated* MLWE *are defined in Figure 1. For an adversary $\mathcal{A}$ trying to solve the* Trunc-MLWE *problem, the respective advantage is defined as*

$$\mathsf{Adv}^{\mathsf{S\text{-}Trunc\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}},c}(\mathcal{A}) = \Pr[\mathsf{S\text{-}Trunc\text{-}MLWE}^{c}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}(\mathcal{A}) = 1],$$

and

$$\mathsf{Adv}^{\mathsf{D\text{-}Trunc\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}},c}(\mathcal{A}) = \Pr[\mathsf{D\text{-}Trunc\text{-}MLWE}^{c}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}(\mathcal{A}) = 1] - \frac{1}{2}.$$

$\underline{\mathsf{S\text{-}Trunc\text{-}MLWE}^{c}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}(\mathcal{A})}$

$\mathbf{U} \leftarrow R_q^{m \times n}$
$\mathbf{A} = \mathsf{Trunc}(\mathbf{U}, c)$
$\mathbf{s} \leftarrow \mathsf{D}_{\mathsf{sec}}^{n}$
$\mathbf{e} \leftarrow \mathsf{D}_{\mathsf{err}}^{m}$
$\mathbf{b} := \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$
$\mathbf{s}' \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b})$
return $\mathbf{s} = \mathbf{s}'$

$\underline{\mathsf{D\text{-}Trunc\text{-}MLWE}^{c}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}(\mathcal{A})}$

$\mathbf{U} \leftarrow R_q^{m \times n}$
$\mathbf{A} = \mathsf{Trunc}(\mathbf{U}, c)$
$b \leftarrow \{0, 1\}$
if $b = 0$:
  $\mathbf{s} \leftarrow \mathsf{D}_{\mathsf{sec}}^{n}$
  $\mathbf{e} \leftarrow \mathsf{D}_{\mathsf{err}}^{m}$
  $\mathbf{b} := \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$
else :
  $\mathbf{b} \leftarrow R_q^{m}$
$b' \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b})$
return $b = b'$

Fig. 1: The experiments for S-Trunc-MLWE and D-Trunc-MLWE.

If no bits are truncated from the matrix, i.e., $c = 0$, we recover the standard MLWE problem [LS15] and simply write $\mathsf{S\text{-}MLWE}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}$ and respectively $\mathsf{D\text{-}MLWE}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}$. If additionally the ring is of degree 1, i.e., $R = \mathbb{Z}$, we recover the standard LWE problem [Reg05].

3.2 Module Learning with Errors with Hints

In the following, we define a variant of the Module Learning with Errors problem, where some approximate hints on the MLWE secret $\mathbf{s}$ are additionally given to the adversary, denoted by Hint-MLWE. Concretely, a hint is of the form $(\mathbf{H}, \mathbf{h})$, with $\mathbf{h} = \mathbf{H}\mathbf{s} + \mathbf{f}$, where $\mathbf{H}$ is the *hint matrix* of bounded infinity norm and $\mathbf{f}$ the *hint noise* coming from some distribution $\mathsf{D}_{\mathsf{noi}}$ over R. Note that $\mathbf{h}$ is in general not taken modulo q.

Different variants of MLWE with hints have been proposed in the literature before. Sometimes, $\mathbf{H}$ is honestly sampled from some distribution [KLSS23].

Other times, it is chosen by the adversary. Here, the adversary either has to choose before seeing the MLWE matrix $\mathbf{A}$, or after having seen it as in [PS24]. In our case, we require the latter case, as we later in Section 5 use Hint-MLWE, where the hint matrix stores the low-order bits of $\mathbf{A}$.

Definition 7 (MLWE with Hints Problem). *Let R be a degree-d ring of integers. Let q, m, n, B and ℓ be positive integers. Further let $\mathsf{D}_{\mathsf{err}}, \mathsf{D}_{\mathsf{sec}}$ and $\mathsf{D}_{\mathsf{noi}}$ be distributions over R. The (interactive) experiments of the search and decision versions of* MLWE *with hints are defined in Figure 2. For an adversary $\mathcal{A}$ trying to solve the* Hint-MLWE *problem, the respective advantage is defined as*

$$\mathsf{Adv}^{\mathsf{S\text{-}Hint\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}},\mathsf{D}_{\mathsf{noi}},B,\ell}(\mathcal{A}) = \Pr[\mathsf{S\text{-}Hint\text{-}MLWE}^{\mathsf{D}_{\mathsf{noi}},B,\ell}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}(\mathcal{A}) = 1],$$

and

$$\mathsf{Adv}^{\mathsf{D\text{-}Hint\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}},\mathsf{D}_{\mathsf{noi}},B,\ell}(\mathcal{A}) = \Pr[\mathsf{D\text{-}Hint\text{-}MLWE}^{\mathsf{D}_{\mathsf{noi}},B,\ell}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}(\mathcal{A}) = 1] - \frac{1}{2}.$$

$\underline{\mathsf{S\text{-}Hint\text{-}MLWE}^{\mathsf{D}_{\mathsf{noi}},B,\ell}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}(\mathcal{A})}$

$\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$
$\mathbb{Z}^{\ell \times n} \ni \mathbf{H} \leftarrow \mathcal{A}(\mathbf{A})$
if $\|\mathbf{H}\|_\infty > B$
  return $\perp$
$\mathbf{s} \leftarrow \mathsf{D}_{\mathsf{sec}}^n$
$\mathbf{e} \leftarrow \mathsf{D}_{\mathsf{err}}^m$
$\mathbf{b} := \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$
$\mathbf{f} \leftarrow \mathsf{D}_{\mathsf{noi}}^\ell$
$\mathbf{h} := \mathbf{H}\mathbf{s} + \mathbf{f} \bmod q$
$\mathbf{s}' \leftarrow \mathcal{A}(\mathbf{b}, \mathbf{h})$
return $\mathbf{s} = \mathbf{s}'$

$\underline{\mathsf{D\text{-}Hint\text{-}MLWE}^{\mathsf{D}_{\mathsf{noi}},B,\ell}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}(\mathcal{A})}$

$\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$
$\mathbb{Z}^{\ell \times n} \ni \mathbf{H} \leftarrow \mathcal{A}(\mathbf{A})$
if $\|\mathbf{H}\|_\infty > B$
  return $\perp$
$\mathbf{s} \leftarrow \mathsf{D}_{\mathsf{sec}}^n$
$b \leftarrow \{0, 1\}$
if $b = 0$:
  $\mathbf{e} \leftarrow \mathsf{D}_{\mathsf{err}}^m$
  $\mathbf{b} := \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$
else :
  $\mathbf{b} \leftarrow \mathbb{Z}_q^m$
$\mathbf{f} \leftarrow \mathsf{D}_{\mathsf{noi}}^\ell$
$\mathbf{h} := \mathbf{H}\mathbf{s} + \mathbf{f} \bmod q$
$b' \leftarrow \mathcal{A}(\mathbf{b}, \mathbf{h})$
return $b = b'$

Fig. 2: The experiments for S-Hint-MLWE and D-Hint-MLWE.

The hardness of Hint-MLWE can be reduced from the hardness of the standard MLWE problem in the case of discrete Gaussian secret and hint noise distributions. This has been proven in the case of RLWE (that is, the special case of MLWE with rank $n = 1$). Our result generalizes their proof to higher ranks, whereas the proof strategy closely follows their original proof. Note that as in the

original proof, we restrict the result to power-of-two cyclotomics. This is mainly due to the fact that we bound the infinity norm of $\mathsf{Rot}(\mathbf{H})$ through the infinity norm of the hint matrix $\mathbf{H}$. One could generalize it to other fields and rings, at the expense of looser norm bounds, determined by the so-called expansion factor [LM06,RSW18].

Theorem 1 (Adapted from [MKMS22, Thm. 1]). *Let R be a power-of-two cyclotomic ring of degree d. Let q, m, n, B and ℓ be positive integers and σ, δ be positive reals such that $\sigma\sqrt{1-\sigma^2B^2d^2n(\ell+2)/\delta^2} \geq \sqrt{d(n+\ell)}$. By $\mathsf{D}_{\mathsf{err}}$ we denote an arbitrary distribution over R. We set $\mathsf{D}_{\mathsf{sec}} = D_{R,\sigma^2}$, $\mathsf{D}'_{\mathsf{sec}} = D_{R,\delta^2}$ and $\mathsf{D}_{\mathsf{noi}} = D_{R,\delta^2}$. Then, there is a reduction from the problem $\mathsf{MLWE}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}$ to $\mathsf{Hint\text{-}MLWE}^{\mathsf{D}_{\mathsf{noi}},B,\ell}_{q,m,n,\mathsf{D}'_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}$. More concretely, assuming that there exists an adversary $\mathcal{A}$ against $\mathsf{Hint\text{-}MLWE}$ with advantage Adv, we can construct an adversary $\mathcal{B}$ against MLWE with advantage at least Adv. The reduction works for both, the search and the decision variants of the problem.*

As example parameters, we can set $\delta^2 \geq 2\sigma^2B^2d^2n(\ell+2)$ and $\sigma \geq \sqrt{2d(n+\ell)}$. When setting $n = 1$, we recover the original result [MKMS22, Thm. 1]. When setting $d = 1$, we obtain the result for plain LWE.

Proof. We detail out the proof in the case of the corresponding decision variants. The proof for the search variants works analogously. Let $\mathcal{A}$ be an adversary against $\mathsf{D\text{-}Hint\text{-}MLWE}$ with advantage Adv. We now construct a reduction $\mathcal{B}$ against $\mathsf{D\text{-}MLWE}$ with advantage at least Adv.

In the $\mathsf{D\text{-}MLWE}$ experiment, $\mathcal{A}$ is given as input $(\mathbf{A}, \mathbf{b}) \in R_q^{m\times n} \times R_q^m$. The reduction now forwards $\mathbf{A}$ as input to the adversary $\mathcal{A}$, who responds with a hint matrix $\mathbf{H} \in R^{\ell\times n}$ such that $\|\mathbf{H}\|_\infty := \|\tau(\mathbf{H})\|_\infty \leq B$.

The matrix $\mathbf{H}$ defines the matrix $\mathsf{Rot}(\mathbf{H}) \in \mathbb{Z}^{d\ell\times dn}$, where every coefficient in $\mathbf{H}$ is replaced by its corresponding multiplication matrix. The matrix $\mathsf{Rot}(\mathbf{H})$ defines a lattice Λ in $R^{d(n+\ell)}$, given as $\Lambda = \left\{(\tau(\mathbf{s}), -\mathsf{Rot}(\mathbf{H})\tau(\mathbf{s}))^T \mid \mathbf{s} \in R^n\right\}$, where τ denotes the coefficient embedding. Every element in Λ can be written as $\mathbf{L}\cdot\tau(\mathbf{s})$ with $\mathbf{L} = (\mathbf{I}_{dn}, -\mathsf{Rot}(\mathbf{H}))^T \in \mathbb{Z}^{d(n+\ell)\times dn}$. By the properties of discrete Gaussians, if $\mathbf{s} \leftarrow \mathsf{D}^n_{\mathsf{sec}} = D_{R^n,\sigma^2}$ (i.e., $\tau(\mathbf{s}) \leftarrow D_{\mathbb{Z}^{dn},\sigma^2}$), then $\mathbf{L}\cdot\tau(\mathbf{s}) \sim D_{\Lambda,\boldsymbol{\Sigma}}$, where $\boldsymbol{\Sigma} = \sigma^2\mathbf{L}\mathbf{L}^T \in \mathbb{Z}^{d(n+\ell)\times d(n+\ell)}$. We further set $\boldsymbol{\Sigma}' = \delta^2\mathbf{I}_{d(n+\ell)} - \boldsymbol{\Sigma}$. As we later use $\boldsymbol{\Sigma}'$ as the covariance matrix of a discrete Gaussian distribution, we have to make sure that it is positive semi-definite.

Claim. Assume $\delta \geq \sigma Bd\sqrt{n(\ell+2)}$, then the matrix $\boldsymbol{\Sigma}'$ is positive semi-definite.

Proof (Claim). Note that $\boldsymbol{\Sigma}$ is symmetric over $\mathbb{Z}$, thus $\boldsymbol{\Sigma}' = \delta^2\mathbf{I}_{n+\ell} - \boldsymbol{\Sigma}$ is symmetric over $\mathbb{Z}$ as well. Let us recall the concrete forms of $\boldsymbol{\Sigma}$ and $\boldsymbol{\Sigma}'$ given by

$$\boldsymbol{\Sigma} = \sigma^2\begin{pmatrix}\mathbf{I}_{dn} & -\mathsf{Rot}(\mathbf{H})^T\\ -\mathsf{Rot}(\mathbf{H}) & \mathsf{Rot}(\mathbf{H})\mathsf{Rot}(\mathbf{H})^T\end{pmatrix},$$

and

$$\boldsymbol{\Sigma}' = \begin{pmatrix}(\delta^2-\sigma^2)\mathbf{I}_{dn} & \sigma^2\mathsf{Rot}(\mathbf{H})^T\\ \sigma^2\mathsf{Rot}(\mathbf{H}) & \delta^2\mathbf{I}_{d\ell} - \sigma^2\mathsf{Rot}(\mathbf{H})\mathsf{Rot}(\mathbf{H})^T\end{pmatrix}.$$

The first dn diagonals of $\boldsymbol{\Sigma}'$ are given by $\delta^2 - \sigma^2$, the last $d\ell$ diagonals are given by $\delta^2 - \sigma^2\|\mathbf{h}_i\|^2$, where $\mathbf{h}_i$ is the i-th row of $\mathsf{Rot}(\mathbf{H})$ for $i \in \{1, \ldots, d\ell\}$. Note that $\|\mathbf{h}_i\|^2 \leq ndB^2$, as we are working over a power-of-two cyclotomic ring. Thus, assuming

$$\delta \geq \sigma B\sqrt{nd}, \tag{2}$$

the diagonal entries of $\boldsymbol{\Sigma}'$ are non-negative. In this case, it is enough to show that $\boldsymbol{\Sigma}'$ is diagonally dominant. By construction, $\|\boldsymbol{\Sigma}\|_\infty \leq \sigma^2 ndB^2$. We write $\boldsymbol{\Sigma}' = (\boldsymbol{\Sigma}'_{ij})_{ij}$ with $i, j \in \{1, \ldots, d(n+\ell)\}$. On the one side, the absolute values of the entries on the diagonal can be lower bounded as $|\boldsymbol{\Sigma}'_{ii}| \geq \delta^2 - \sigma^2 ndB^2$ for $i \in \{1, \ldots, d(n+\ell)\}$. On the other side, the sum of the absolute values of the entries off the diagonal can be upper bounded as

$$\sum_{j \neq i} \left|\boldsymbol{\Sigma}'_{ij}\right| \leq \max\left\{\sigma^2 dnB + \sigma^2 d^2 \ell n B^2, \sigma^2 d\ell B\right\} \leq \sigma^2 B^2 d^2 n(\ell+1),$$

where we used that B, d, ℓ are positive integers. Overall $\boldsymbol{\Sigma}'$ is diagonally dominant if

$$\delta \geq \sigma B d\sqrt{n(\ell+2)}. \tag{3}$$

■

The reduction $\mathcal{B}$ continues as follows. They sample $(\mathbf{s}', \mathbf{f}') \leftarrow D_{R^{n+\ell}, \boldsymbol{\Sigma}'}$ (i.e., $(\tau(\mathbf{s}'), \tau(\mathbf{f}')) \leftarrow D_{\mathbb{Z}^{d(n+\ell)}, \boldsymbol{\Sigma}'}$) and set $\mathbf{b}' = \mathbf{b} + \mathbf{A}\mathbf{s}'$ and $\mathbf{h} = \mathbf{H}\mathbf{s}' + \mathbf{f}'$. They then forward $(\mathbf{b}', \mathbf{h})$ to the adversary $\mathcal{A}$. On the output bit b by $\mathcal{A}$, the reduction also outputs b as their answer. We now analyze the advantage of $\mathcal{B}$, assuming that $\mathcal{A}$ has advantage Adv.

Case 1) Assume that $(\mathbf{A}, \mathbf{b})$ is given as $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$. Thus, $\mathbf{b}' = \mathbf{A}(\mathbf{s} + \mathbf{s}') + \mathbf{e}$. Set $\mathbf{f} = -\mathbf{H}\mathbf{s} + \mathbf{f}'$. Note that the values $\mathbf{s}$ and $\mathbf{f}$ are not known to the reduction $\mathcal{B}$, but only needed to argue that $(\mathbf{b}', \mathbf{h})$ has the right distribution. Then, $(\mathbf{s}+\mathbf{s}', \mathbf{f}) = (\mathbf{s} + \mathbf{s}', -\mathbf{H}\mathbf{s} + \mathbf{f}') = (\mathbf{s}, -\mathbf{H}\mathbf{s}) + (\mathbf{s}', \mathbf{f}')$, with $(\mathbf{s}, -\mathbf{H}\mathbf{s}) \sim D_{R^{n+\ell}, \boldsymbol{\Sigma}}$ and $(\mathbf{s}', \mathbf{f}') \sim D_{R^{n+\ell}, \boldsymbol{\Sigma}'}$. By Lemma 3 Item 2, this implies that $(\mathbf{s}+\mathbf{s}', \mathbf{f}) \sim D_{R^{n+\ell}, \delta^2}$ as long as the eigenvalues of $\boldsymbol{\Gamma} = \sigma\sqrt{\mathbf{I}_{d(n+\ell)} - \sigma^2 \mathbf{L}\mathbf{L}^T/\delta^2}$ are above the smoothing parameter of $\mathbb{Z}^{d(n+\ell)}$. To lower bound the eigenvalues of $\boldsymbol{\Gamma}$, it suffices to upper bound the eigenvalues of $\mathbf{L}\mathbf{L}^T$. To do so, we use a known result of spectral theory: It states that the eigenvalues of $\mathbf{L}\mathbf{L}^T$ can be upper bounded by the sum of the absolute values of any of its row. The latter can be upper bounded by $B^2 d^2 n(\ell+2)$, implying the condition

$$\sigma\sqrt{1 - \sigma^2 B^2 d^2 n(\ell+2)/\delta^2} \geq \eta(\mathbb{Z}^{d(n+\ell)}). \tag{4}$$

Note that Equation 4 subsumes Equations 3 and 2. Moreover, $\eta(\mathbb{Z}^{d(n+\ell)}) \leq \sqrt{d(n+\ell)}$ and thus the condition of the theorem statement fulfills the above conditions. It yields $\mathbf{H}(\mathbf{s}+\mathbf{s}')+\mathbf{f} = \mathbf{H}(\mathbf{s}+\mathbf{s}')-\mathbf{H}\mathbf{s}+\mathbf{f}' = \mathbf{H}\mathbf{s}'+\mathbf{f}'$ and thus $(\mathbf{b}', \mathbf{h})$ is distributed correctly.

Case 2) Assuming $(\mathbf{A}, \mathbf{b})$ comes from the uniform distribution, so does $(\mathbf{A}, \mathbf{b}')$. With the same argumentation as above, $\mathbf{h}$ has the correct distribution and hence $(\mathbf{b}', \mathbf{h})$ is a valid input to $\mathcal{A}$, concluding the proof. □

4 Hardness of Truncated MLWE Using Rényi Divergence

We begin with our first approach to reduce the hardness of Trunc-MLWE from standard MLWE, using the Rényi divergence. Note that the results only apply to the respective search variants.

Theorem 2. *Let R be a power-of-two cyclotomic ring of degree d. Further, let α, q, m, n, η and c be positive integers and δ be a positive real. Further, let $\mathsf{D}_{\mathsf{sec}}, \mathsf{D}_{\mathsf{err}}$ and $\mathsf{D}'_{\mathsf{err}}$ be distributions over R such that $\mathsf{D}_{\mathsf{sec}}$ is (η, δ)-bounded (Def. 1). Then, there is a reduction from the problem $\mathsf{S\text{-}MLWE}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}'_{\mathsf{err}}}$ to the problem $\mathsf{S\text{-}Trunc\text{-}MLWE}^{c}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}$. More concretely, assuming there exists an adversary $\mathcal{A}$ against S-Trunc-MLWE with advantage Adv, we can transform them into an adversary against S-MLWE with advantage Adv′ such that*

$$(\mathsf{Adv} - \delta^n)^{\frac{\alpha}{\alpha-1}} \leq \mathsf{Adv}' \cdot \mathsf{RD}_\alpha(\mathsf{D}_{\mathsf{err}} + \vec{\mu} \| \mathsf{D}'_{\mathsf{err}})^m,$$

where $\vec{\mu} = (\mu, \ldots, \mu) \in \mathbb{Z}^d \cong R$, with $\mu = d \cdot 2^c \cdot n \cdot \eta$.

The reduction loss is reflected by replacing the error distribution $\mathsf{D}_{\mathsf{err}}$ by a wider distribution $\mathsf{D}'_{\mathsf{err}}$. How much wider the distribution has to be is impacted by the ring degree d, the number of truncated bits c, the bound on secrets η and the rank n of the MLWE problem.

Proof. Let $\mathcal{A}$ be an adversary against S-Trunc-MLWE, whose experiment is defined in Figure 1. On input $(\mathbf{A}, \mathbf{b})$, the adversary $\mathcal{A}$ outputs a guess $\mathbf{s}'$ and wins the experiment if the guess was correct. Below, we argue the theorem via a series of intermediate hybrids, each specifying the distribution of the input $(\mathbf{A}, \mathbf{b})$ given to $\mathcal{A}$.

H_0: Sample $(\mathbf{A}, \mathbf{b})$ as specified in the S-Trunc-MLWE game in Figure 1. Thus, $\mathbf{A} = \mathsf{Trunc}(\mathbf{U}, c)$ with $\mathbf{U} \leftarrow R_q^{m \times n}$. Moreover, $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$, with $\mathbf{s} \leftarrow \mathsf{D}_{\mathsf{sec}}^n$ and $\mathbf{e} \leftarrow \mathsf{D}_{\mathsf{err}}^m$.

H_1: We now sample $\mathbf{A}$ as in H_0, but change how $\mathbf{b}$ is defined. First, we sample $\mathbf{f} \leftarrow (\mathsf{D}'_{\mathsf{err}})^m$ and then we set $\mathbf{b} = \mathbf{U}\mathbf{s} + \mathbf{f} \bmod q$.

H_2: Now, $\mathbf{b}$ is sampled as in H_1. However, we modify the input $\mathbf{A}$ by setting it to $\mathbf{A} := \mathbf{U}$. In other words, $(\mathbf{A}, \mathbf{b})$ corresponds to an instance of S-MLWE (without truncation).

From H_0 to H_1: First, we use that $\mathsf{D}_{\mathsf{sec}}$ is (η, δ)-bounded, to condition on the event of $\|\mathbf{s}\|_\infty \leq \eta$ implying

$$\begin{aligned}
\mathsf{Adv} &= \Pr[\mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b}) \mid (\mathbf{A}, \mathbf{b}) \sim \mathsf{H}_0] \\
&= \Pr[\mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b}) \mid (\mathbf{A}, \mathbf{b}) \sim \mathsf{H}_0 \wedge \|\mathbf{s}\|_\infty \leq \eta] \cdot \Pr[\|\mathbf{s}\|_\infty \leq \eta] \\
&\quad + \Pr[\mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b}) \mid (\mathbf{A}, \mathbf{b}) \sim \mathsf{H}_0 \wedge \|\mathbf{s}\|_\infty > \eta] \cdot \Pr[\|\mathbf{s}\|_\infty > \eta] \\
&\leq \Pr[\mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b}) \mid (\mathbf{A}, \mathbf{b}) \sim \mathsf{H}_0 \wedge \|\mathbf{s}\|_\infty \leq \eta] + \delta^n.
\end{aligned}$$

Second, we use the probability preservation property of the Rényi divergence from Lemma 5 to argue

$$\begin{aligned}&\Pr[\mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b}) \mid (\mathbf{A}, \mathbf{b}) \sim \mathsf{H}_0 \wedge \|\mathbf{s}\|_\infty \leq \eta]^{\frac{\alpha}{\alpha-1}} \\ \leq\ &\Pr[\mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b}) \mid (\mathbf{A}, \mathbf{b}) \sim \mathsf{H}_1 \wedge \|\mathbf{s}\|_\infty \leq \eta] \cdot \mathsf{RD}_\alpha(\mathsf{H}_0 \| \mathsf{H}_1).\end{aligned}$$

Note that the only difference between the two hybrids is how $\mathbf{b}$ is defined. In H_0, $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} = (\mathbf{U} - \mathbf{N}_\mathbf{U})\mathbf{s} + \mathbf{e} = \mathbf{U}\mathbf{s} + (\mathbf{e} - \mathbf{N}_\mathbf{U}\mathbf{s})$, where $\mathbf{N}_\mathbf{U} = \mathbf{U} \bmod 2^c$. In H_1, $\mathbf{b} = \mathbf{U}\mathbf{s} + \mathbf{f}$. By the data processing inequality from Lemma 6, consider the function $f(y) = (\mathbf{A}, \mathbf{A}\mathbf{s} + y)$, where $\mathbf{N}_\mathbf{U}\mathbf{s}$ serves as a fixed parameter for the distributions and the function; however, only the first distribution uses $\mathbf{N}_\mathbf{U}\mathbf{s}$, subtracting it from every sample. This yields:

$$\begin{aligned}\mathsf{RD}_\alpha(\mathsf{H}_0 \| \mathsf{H}_1) &\leq \mathsf{RD}_\alpha(\mathbf{e} - \mathbf{N}_\mathbf{U}\mathbf{s} \| \mathbf{f}) \\ &\leq \mathsf{RD}_\alpha(\mathsf{D}_{\mathsf{err}} + \vec{\mu} \| \mathsf{D}'_{\mathsf{err}})^m.\end{aligned}$$

where $\|-\mathbf{N}_\mathbf{U}\mathbf{s}\|_\infty \leq d \cdot 2^c \cdot n \cdot \eta = \mu$. Note that here we are using the properties of power-of-two cyclotomics. When generalizing to other rings, the bound would be looser by the so-called expansion factor [LM06,RSW18]. Overall,

$$(\mathsf{Adv} - \delta^n)^{\frac{\alpha}{\alpha-1}} \leq \Pr[\mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b}) \mid (\mathbf{A}, \mathbf{b}) \sim \mathsf{H}_1] \cdot \mathsf{RD}_\alpha(\mathsf{D}_{\mathsf{err}} + \vec{\mu} \| \mathsf{D}'_{\mathsf{err}})^m.$$

From H_1 to H_2: We observe that the only difference between H_1 and H_2 is that the c least significant bits of $\mathbf{U}$ are removed in H_1, but not in H_2. Removing the low-order bits only results in less information being transmitted, so the advantage of an adversary in H_1 cannot be greater than that of an adversary in H_2:

$$\Pr[\mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b}) \mid (\mathbf{A}, \mathbf{b}) \sim \mathsf{H}_1] \leq \Pr[\mathbf{s} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{b}) \mid (\mathbf{A}, \mathbf{b}) \sim \mathsf{H}_2].$$

The proof concludes by observing that H_2 is equivalent to an instance of the problem $\mathsf{S\text{-}MLWE}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}'_{\mathsf{err}}}$. □

Theorem 2 applies to any secret and error distributions as long as the secrets are of bounded infinity norm and the Rényi divergences are well-defined and small enough. We now provide example corollaries for bounded uniform, rounded and discrete Gaussian distributions.

Corollary 1. *Let R be a power-of-two cyclotomic ring of degree d. Further, let α, q, m, n, η and c be positive integers. We set $\mathsf{D}_{\mathsf{sec}} = \mathbf{U}_{0,\eta+1}$ and $\mathsf{D}_{\mathsf{err}} = \mathbf{U}_{0,k}$ and $\mathsf{D}'_{\mathsf{err}} = \mathbf{U}_{0,k+\mu}$, where $\mu = d \cdot n \cdot 2^c \cdot \eta$. Then, for any adversary $\mathcal{A}$ it holds*

$$\mathsf{Adv}^{\mathsf{S\text{-}Trunc\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}},c}(\mathcal{A})^{\frac{\alpha}{\alpha-1}} \leq \mathsf{Adv}^{\mathsf{S\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}'_{\mathsf{err}}}(\mathcal{A}) \cdot \left(\frac{k+\mu}{k}\right)^{d \cdot m}.$$

Here, it is very clear that the resulting error distribution $\mathsf{D}'_{\mathsf{err}}$ has infinity norm larger than 2^c, avoiding the trivial setup mentioned in the introduction.

Remark 1. In order to bound the reduction loss, one has to make sure that k is large enough. For instance, when $k \geq d^2mn$, the loss is bounded above by

$$\left(\frac{k+\mu}{k}\right)^{d\cdot m} = \left(1+\frac{\mu}{k}\right)^{d\cdot m} \leq \left(1+\frac{\eta\cdot 2^c}{dm}\right)^{dm} \leq e^{\eta\cdot 2^c},$$

which is constant for η and c being constants.

Proof. We observe that $\mathsf{D}_{\mathsf{sec}} = \mathbf{U}_{0,\eta+1}$ is $(\eta, 0)$-bounded. We now provide a concrete value for the Rényi divergence from Theorem 2. We observe that $\mathsf{D}_{\mathsf{err}} + \vec{\mu} = \mathbf{U}_{0,k} + \vec{\mu} = \mathbf{U}_{\mu,k+\mu}$ and $\mathsf{Supp}(\mathsf{D}_{\mathsf{err}} + \vec{\mu}) = \mathsf{Supp}(\mathbf{U}_{\mu,k+\mu}) = \{\mu, \ldots, \mu+k-1\}^d \subset \{0, \ldots, \mu+k-1\}^d = \mathsf{Supp}(\mathbf{U}_{0,k+\mu}) = \mathsf{Supp}(\mathsf{D}'_{\mathsf{err}})$, implying a well-defined Rényi divergence.

$$\begin{aligned}\mathsf{RD}_\alpha(\mathsf{D}_{\mathsf{err}} + \vec{\mu} \| \mathsf{D}'_{\mathsf{err}}) &= \left(\sum_{x\in\mathsf{Supp}(\mathbf{U}_{\mu,k+\mu})} \frac{\mathbf{U}_{\mu,k+\mu}(x)^\alpha}{\mathbf{U}_{0,k+\mu}(x)^{\alpha-1}}\right)^{1/(\alpha-1)} \\ &= \left((k)^d \cdot \frac{(1/(k))^{d\alpha}}{(1/(k+\mu))^{d(\alpha-1)}}\right)^{1/(\alpha-1)} = \left(\frac{k+\mu}{k}\right)^d.\end{aligned}$$

□

Corollary 2. *Let R be a power-of-two cyclotomic ring of degree d. Further, let α, q, m, n, η and c be positive integers and $\mathsf{D}_{\mathsf{sec}}$ be any (η, δ)-bounded distribution over R. We set $\mathsf{D}_{\mathsf{err}} = \mathsf{D}'_{\mathsf{err}} = \lfloor D_{\sigma^2} \rceil$ over R (via coefficient embedding τ) for positive real σ. Then, for any adversary $\mathcal{A}$ it holds*

$$\left(\mathsf{Adv}^{\mathsf{S\text{-}Trunc\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}},c}(\mathcal{A}) - \delta^n\right)^{\frac{\alpha}{\alpha-1}} \leq \mathsf{Adv}^{\mathsf{S\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}'_{\mathsf{err}}}(\mathcal{A}) \cdot \exp\left(\frac{\alpha dm\mu^2}{2\sigma^2}\right),$$

where $\mu = d \cdot n \cdot 2^c \cdot \eta$.

Proof. We provide a concrete value for the Rényi divergence in Theorem 2. We use Lemma 6 with $f(\mathbf{x}) = \lfloor \mathbf{x} \rceil$ and $\vec{\mu} \in \mathbb{Z}^d$ to argue that

$$\mathsf{RD}_\alpha(\mathsf{D}_{\mathsf{err}} + \vec{\mu} \| \mathsf{D}'_{\mathsf{err}}) = \mathsf{RD}_\alpha(\lfloor D_{\mu,\sigma^2} \rceil \| \lfloor D_{\sigma^2} \rceil) \leq \mathsf{RD}_\alpha(D_{\mu,\sigma^2} \| D_{\sigma^2}).$$

By Lemma 7 it yields

$$\mathsf{RD}_\alpha(D_{\mu,\sigma^2} \| D_{\sigma^2}) \leq \exp\left(\frac{\alpha d\mu^2}{2\sigma^2}\right).$$

Thus, overall

$$\mathsf{RD}_\alpha(\mathsf{D}_{\mathsf{err}} + \mu \| \mathsf{D}'_{\mathsf{err}})^m = \exp\left(\frac{m\alpha d\mu^2}{2\sigma^2}\right),$$

concluding the proof. □

Corollary 3. *Let R be a power-of-two cyclotomic ring of degree d. Further, let α, q, m, n, η and c be positive integers and $\mathsf{D}_{\mathsf{sec}}$ be any (η, δ)-bounded distribution over R. We set $\mathsf{D}_{\mathsf{err}} = \mathsf{D}'_{\mathsf{err}} = D_{R,\sigma^2}$ for positive real σ. Then, for any adversary $\mathcal{A}$ it holds*

$$\left(\mathsf{Adv}^{\mathsf{S\text{-}Trunc\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}},c}(\mathcal{A}) - \delta^n\right)^{\frac{\alpha}{\alpha-1}} \leq \mathsf{Adv}^{\mathsf{S\text{-}MLWE}}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}'_{\mathsf{err}}}(\mathcal{A}) \cdot \exp\left(\frac{\alpha d m \mu^2}{2\sigma^2}\right),$$

where $\mu = d \cdot n \cdot 2^c \cdot \eta$.

Proof. As before, we provide concrete values for the two corresponding Rényi divergences in Theorem 2. We observe that

$$\mathsf{RD}_\alpha(\mathsf{D}_{\mathsf{err}} + \vec{\mu} \| \mathsf{D}'_{\mathsf{err}})^m = \mathsf{RD}_\alpha(D_{R,\vec{\mu},\sigma^2} \| D_{R,\sigma^2})^m = \mathsf{RD}_\alpha(D_{R^m,\boldsymbol{\mu},\sigma^2} \| D_{R^m,\sigma^2}).$$

where $\boldsymbol{\mu} = (\vec{\mu}, \ldots, \vec{\mu})^T \in R^m$. By applying Lemma 8 and noticing that $\boldsymbol{\mu} \in \mathbb{Z}^{md}$ and therefore $\epsilon = 0$, we can argue that

$$\mathsf{RD}_\alpha(D_{R^m,\boldsymbol{\mu},\sigma^2} \| D_{R^m,\sigma^2}) \leq \left(\frac{1+\epsilon}{1-\epsilon}\right)^{\alpha/\alpha-1} \cdot e^{\alpha \|\boldsymbol{\mu}\|^2/(2\sigma^2)} = e^{\alpha d m \mu^2/(2\sigma^2)},$$

with $\|\boldsymbol{\mu}\| = md\mu^2$, concluding the proof. □

4.1 Public Sampleability Does Not Help Here

As of today, the only way to use the Rényi divergence for decision problems is to use the public sampleability framework from [BLR+18, Sec. 4]. We argue below that, even though the truncated MLWE problem can be expressed as a publicly sampleable problem, it does not lead to meaningful results. Intuitively, the problem is that truncated and uniform matrices are not Rényi close, which would be needed.

More formally, we observe that truncated MLWE fits well into the public sampleable setting: given $(\mathbf{A}, \mathbf{b})$, one can easily sample fresh samples by drawing new MLWE secret and errors. A very similar setting was used in [BLR+18, Sec. 4.2]. However, the loss in the advantage depends on the Rényi divergence between the standard way to sample $\mathbf{A}$ (i.e. uniformly over R_q) and the truncated way to sample it (i.e. cutting the c lowest order bits off). Unfortunately, their Rényi divergence is given by 2^{cdmn}, which is exponential in the product of ring degree and matrix dimensions, which is usually linear proportional to the security parameter. Given this big loss, the reduction becomes vacuous.

5 Hardness of Truncated Module LWE Using Hints

We continue with our second approach to reduce the hardness of Trunc-MLWE from standard MLWE, going through the intermediate Hint-MLWE problem (cf. Section 3.2). The result now applies to both search and decision variants, but requires a decomposition property for the error/noise distributions.

The following theorem establishes a reduction from Hint-MLWE to Trunc-MLWE.

Theorem 3. *Let R be a ring of integers of degree d, and let q, m, n, B and c be positive integers such that $B = 2^c$. Moreover, let $\mathsf{D}_{\mathsf{sec}}, \mathsf{D}_{\mathsf{err}}, \mathsf{D}'_{\mathsf{err}}$ and $\mathsf{D}_{\mathsf{noi}}$ be distributions over R such that $(\mathsf{D}'_{\mathsf{err}})^m$ is statistically close to $(\mathsf{D}_{\mathsf{err}})^m + (\mathsf{D}_{\mathsf{noi}})^m$. Then, there is a reduction from the problem $\mathsf{Hint\text{-}MLWE}^{\mathsf{D}_{\mathsf{noi}},B,m}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}$ to the problem $\mathsf{Trunc\text{-}MLWE}^{c}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}'_{\mathsf{err}}}$. More concretely, assuming that there exists an adversary $\mathcal{A}$ against $\mathsf{Trunc\text{-}MLWE}$ with advantage Adv, we can construct an adversary $\mathcal{B}$ against $\mathsf{Hint\text{-}MLWE}$ with advantage at least Adv. The reduction works for both, the search and the decision variants of the problems.*

Proof. We detail out the proof in the case of the corresponding decision variants. The proof for the search variants works analogously. Let $\mathcal{A}$ be an adversary against $\mathsf{Trunc\text{-}MLWE}$ with advantage Adv. We now construct a reduction $\mathcal{B}$ against $\mathsf{Hint\text{-}MLWE}$ with advantage at least Adv.

In the $\mathsf{Hint\text{-}MLWE}$ experiment, $\mathcal{B}$ is given as input $\mathbf{A}$, sampled uniformly at random from $R_q^{m\times n}$. The reduction $\mathcal{B}$ now outputs $\mathbf{H} := -(\mathbf{A} \bmod 2^c) \in R_q^{m\times n}$ as hint matrix. Upon receiving $(\mathbf{b}, \mathbf{h}) \in R_q^m \times R_q^m$, the reduction $\mathcal{B}$ sets $\mathbf{t} := (\mathbf{b} + \mathbf{h})^T \in R_q^m$ as well as $\mathbf{B} := (\mathbf{A} + \mathbf{H})^T \in R_q^{m\times n}$. Then, they provide $(\mathbf{B}, \mathbf{t})$ as input to $\mathcal{A}$. Let b' be the guess of $\mathcal{A}$, then $\mathcal{B}$ forwards b' as their guess.

First, we observe that $\|\mathbf{H}\|_\infty \leq 2^c = B$, hence $\mathbf{H}$ is a valid hint matrix for $\mathsf{Hint\text{-}MLWE}$, as specified in the $\mathsf{Hint\text{-}MLWE}$ experiment in Figure 2. Further, we see that $\mathbf{A} + \mathbf{H} = \mathbf{A} - (\mathbf{A} \bmod 2^c) = \mathsf{Trunc}(\mathbf{A}, c)$. Hence, $\mathbf{B}$ has the correct distribution as specified in the $\mathsf{Trunc\text{-}MLWE}$ experiment in Figure 1.

Case 1: If $b = 0$ in the $\mathsf{Hint\text{-}MLWE}$ experiment, then $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$ and $\mathbf{h} = \mathbf{H}\mathbf{s} + \mathbf{f}''$ for some $\mathbf{s} \leftarrow \mathsf{D}_{\mathsf{sec}}^n$, $\mathbf{e} \leftarrow \mathsf{D}_{\mathsf{err}}^m$ and $\mathbf{f}'' \leftarrow \mathsf{D}_{\mathsf{noi}}^m$. Hence, $\mathbf{t} = \mathbf{A}\mathbf{s} + \mathbf{e} + \mathbf{h} = (\mathbf{A}+\mathbf{H})\mathbf{s}+\mathbf{e}+\mathbf{f} = \mathsf{Trunc}(\mathbf{A}, c)\mathbf{s}+\mathbf{e}+\mathbf{f}$. Note that $(\mathbf{e}+\mathbf{f})^T$ is statistically close to some $\mathbf{g} \leftarrow (\mathsf{D}'_{\mathsf{err}})^m$. Overall, $(\mathbf{B}, \mathbf{t})$ is statistically close to the input distribution in case of $b = 0$ in $\mathsf{Trunc\text{-}MLWE}$.

Case 2: If $b = 1$ in the $\mathsf{Hint\text{-}MLWE}$ experiment, then $\mathbf{b} \leftarrow R_q^m$ and $\mathbf{h}$ and $\mathbf{f}'$ are chosen independently of $\mathbf{b}$. Thus, $\mathbf{t}$ is also distributed uniformly at random over R_q^m and hence $(\mathbf{B}, \mathbf{t})$ corresponds to the input distribution in case of $b = 1$ in $\mathsf{Trunc\text{-}MLWE}$.

If $\mathcal{A}$ succeeds to guess correctly (i.e., $b' = b$) with probability Adv in the experiment of $\mathsf{Trunc\text{-}MLWE}$, then $\mathcal{B}$ succeeds to guess correctly in the game of $\mathsf{Hint\text{-}MLWE}$ with probability at least Adv as well. □

The following corollary instantiates the above theorem with discrete Gaussian distributions over power-of-two cyclotomic rings and combines it with Theorem 1 to provide the complete reduction from MLWE to $\mathsf{Trunc\text{-}MLWE}$.

Corollary 4. *Let R be a power-of-two cyclotomic ring of degree d. Further, let q, m, n and c be positive integers. Moreover, let γ, δ and σ be positive reals such that $\sigma \geq \sqrt{2d(n+m)}$, $\delta \geq \sigma 2^c d\sqrt{2n(m+2)}$, and $\sigma^2 + \delta^2 = \gamma^2$. We set $\mathsf{D}_{\mathsf{sec}} = \mathsf{D}_{\mathsf{err}} = D_{R,\sigma^2}$, $\mathsf{D}'_{\mathsf{sec}} = D_{R,\delta^2}$ and $\mathsf{D}'_{\mathsf{err}} = D_{R,\gamma^2}$. Then, there is a reduction from $\mathsf{MLWE}_{q,m,n,\mathsf{D}_{\mathsf{sec}},\mathsf{D}_{\mathsf{err}}}$ to the problem $\mathsf{Trunc\text{-}MLWE}^{c}_{q,m,n,\mathsf{D}'_{\mathsf{sec}},\mathsf{D}'_{\mathsf{err}}}$.*

Proof. Note that $(D_{R,\sigma^2})^m = D_{R^m,\sigma^2}$ for every positive real σ. The corollary follows by Lemma 3 Item 1 with $\Lambda = R^m$ and noting that $\eta(R^m) \leq \sqrt{dm}$, thus $\sigma, \delta \geq \sqrt{2} \cdot \eta(R^m)$. Moreover, we instantiate Theorem 1 with $\ell = m$ and MLWE noise distributed as D_{R,σ^2}. □

6 Comparison

As explained in the introduction, our work closes an open problem left open by [JZW+23], by providing a reduction from standard MLWE to the truncated problem in its decision variant. Their work only provides a reduction from the less standard module variant of NTRU, denoted by MNTRU, and is limited to the search versions. In total, we describe two approaches. The two different reductions in Theorem 2 and Theorem 3 come with different advantages and disadvantages, as detailed out in the following and summarized in Table 1.

Result	Assumption	Variant	Secret	Error
[JZW+23]	MNTRU	Search	Entropic	Gaussian
Theorem 2	MLWE	Search	Bounded	Rényi-close
Theorem 3	MLWE	Decision	Gaussian	Gaussian

Table 1: Comparison of [JZW+23] with our two results to prove the hardness of Trunc-MLWE. Entropic distributions denote any distribution with enough min-entropy. Rényi-close denotes the fact that the (shifted) starting and ending error distribution have to be Rényi-close.

Note that all reductions, including [JZW+23] and ours, preserve the ring degree d, the modulus q, as well as the MLWE dimensions m and n.

The reduction of [JZW+23] works for any secret distribution which has enough min-entropy, what we denote by an *entropic* distribution. Thus, it also covers secret distributions which have large infinity norm. Both the distribution of the starting MNTRU problem and the final error distribution of truncated MLWE are assumed to be discrete Gaussians.

The most important positive aspect of the reduction in Theorem 2 using the Rényi divergence (Section 4) is its flexibility in terms of considered secret and error distributions. In particular, the reduction preserves the secret distribution $\mathsf{D}_{\mathsf{sec}}$, which can be any (η, δ)-bounded distribution over R. Moreover, it can be instantiated with various error distributions, as long as their corresponding Rényi divergences are defined and small. Of course, only error distributions for which the starting MLWE problem is hard are useful. Corollaries 1, 2 and 3 give three concrete examples for bounded uniform, rounded Gaussian, and discrete Gaussian, but these are by far not the only ones possible. On the other hand, the reduction is limited to the search variants, as the public sampleability result in Section 4.1 is vacuous.

The biggest advantage of the reduction in Theorem 3 using hints (Section 5) is that it works for the decision variant, which is needed when using it in the context of IND-CPA secure public-key encryption schemes, for example to compressed public keys. However, both Theorems 1 and 3 make use of the decomposition theorems of discrete Gaussians. Thus, our overall result in Corollary 4 is limited to discrete Gaussians secret and error distributions. Furthermore, the reduction significantly increases the width of the discrete Gaussian secret distribution.

Finally, we concretely compare the parameter conditions of our results for the case of discrete Gaussians, as summarized in Table 2. For simplicity, we set the starting secret distribution to $\mathsf{D}_{\mathsf{sec}} = D_{R,\beta^2}$ for some positive real β. By Lemma 1, the distribution $\mathsf{D}_{\mathsf{sec}} = D_{R,\beta^2}$ is then $(\sqrt{d}\beta, \mathsf{negl}(d))$-bounded. Then, both Corollary 3 and Corollary 4 reduce $\mathsf{MLWE}_{q,m,n,D_{R,\beta^2},D_{R,\sigma^2}}$ to the problem $\mathsf{Trunc\text{-}MLWE}^{c}_{q,m,D_{R,\delta^2},D_{R,\gamma^2}}$. One can observe that the loss in the ring degree d parameter is larger in the first result (left column) than in the second result (right column). The opposite effect can be observed for the matrix dimension m. Overall, the second reduction sets more constraints than the first. The latter is thus preferable in the case of large degrees and small dimension m, as one observes in practice.

	S-Trunc-MLWE (Cor. 3 & RD_α)	D-Trunc-MLWE (Cor. 4)
β	> 0	$= \sigma$
σ	$\geq \beta d^2 2^c n\sqrt{\alpha m}$	$\geq \sqrt{2d(n+m)}$
δ	$= \beta$	$\geq \sigma 2^c d\sqrt{2n(m+2)}$
γ	$= \sigma$	$= \sqrt{\sigma^2 + \delta^2}$

Table 2: Parameter comparison between the two different reductions from MLWE to Trunc-MLWE for discrete Gaussian secret and error distributions.

Acknowledgement

The research described in this paper has received funding from the European Research Council (ERC) under the European Unions's Horizon 2020 research and innovation programme under grant agreement No 803096 (SPEC) and the Danish Independent Research Council under Grant-ID DFF-2064-00016B (YOSO).

We thank Corentin Jeudy for pointing us to the work of Boneh et al. [BLMR13], which we had missed. Moreover, we thank the anonymous reviewers mentioning the trivial setup to us.

References

ACPS09. Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning

problems. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, Heidelberg, August 2009.

BCD+16. Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! Practical, quantum-secure key exchange from LWE. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1006–1018. ACM Press, October 2016.

BD20. Zvika Brakerski and Nico Döttling. Hardness of LWE on general entropic distributions. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 551–575. Springer, Heidelberg, May 2020.

BJRW20. Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois, and Weiqiang Wen. Towards classical hardness of module-LWE: The linear rank case. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 289–317. Springer, Heidelberg, December 2020.

BJRW22. Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois, and Weiqiang Wen. Entropic hardness of module-lwe from module-ntru. In Takanori Isobe and Santanu Sarkar, editors, *Progress in Cryptology - INDOCRYPT 2022 - 23rd International Conference on Cryptology in India, Kolkata, India, December 11-14, 2022, Proceedings*, volume 13774 of *Lecture Notes in Computer Science*, pages 78–99. Springer, 2022.

BJRW23. Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois, and Weiqiang Wen. On the hardness of module learning with errors with short distributions. *Journal of Cryptology*, 36(1):1, January 2023.

BLL+15. Shi Bai, Adeline Langlois, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 3–24. Springer, Heidelberg, November / December 2015.

BLMR13. Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 410–428. Springer, Heidelberg, August 2013.

BLP+13. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 575–584. ACM Press, June 2013.

BLR+18. Shi Bai, Tancrède Lepoint, Adeline Roux-Langlois, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. *Journal of Cryptology*, 31(2):610–640, April 2018.

CPS+20. Chitchanok Chuengsatiansup, Thomas Prest, Damien Stehlé, Alexandre Wallet, and Keita Xagawa. ModFalcon: Compact signatures based on module-NTRU lattices. In Hung-Min Sun, Shiuh-Pyng Shieh, Guofei Gu, and Giuseppe Ateniese, editors, *ASIACCS 20*, pages 853–866. ACM Press, October 2020.

DM13. Nico Döttling and Jörn Müller-Quade. Lossy codes and a new variant of the learning-with-errors problem. In Thomas Johansson and Phong Q.

Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 18–34. Springer, Heidelberg, May 2013.

dPKPR24. Rafaël del Pino, Shuichi Katsumata, Thomas Prest, and Mélissa Rossi. Raccoon: A masking-friendly signature proven in the probing model. In *CRYPTO (1)*, volume 14920 of *Lecture Notes in Computer Science*, pages 409–444. Springer, 2024.

GKPV10. Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In Andrew Chi-Chih Yao, editor, *ICS 2010*, pages 230–240. Tsinghua University Press, January 2010.

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008.

HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998.

JLS24. Aayush Jain, Huijia Lin, and Sagnik Saha. A systematic study of sparse LWE. In *CRYPTO 2024, Part III*, LNCS, pages 210–245, August 2024.

JZW+23. Wenjuan Jia, Jiang Zhang, Baocang Wang, et al. Hardness of module-lwe with semiuniform seeds from module-ntru. *IET Information Security*, 2023, 2023.

KLSS23. Duhyeong Kim, Dongwon Lee, Jinyeong Seo, and Yongsoo Song. Toward practical lattice-based proof of knowledge from hint-mlwe. In *CRYPTO (5)*, volume 14085 of *Lecture Notes in Computer Science*, pages 549–580. Springer, 2023.

LM06. Vadim Lyubashevsky and Daniele Micciancio. Generalized compact Knapsacks are collision resistant. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP 2006, Part II*, volume 4052 of *LNCS*, pages 144–155. Springer, Heidelberg, July 2006.

LS15. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.*, 75(3):565–599, 2015.

LSS14. Adeline Langlois, Damien Stehlé, and Ron Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 239–256. Springer, Heidelberg, May 2014.

LWZW24. Hao Lin, Mingqiang Wang, Jincheng Zhuang, and Yang Wang. Hardness of entropic module-lwe. *Theor. Comput. Sci.*, 999:114553, 2024.

Lyu12. Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Heidelberg, April 2012.

Mic18. Daniele Micciancio. On the hardness of learning with errors with binary secrets. *Theory Comput.*, 14(1):1–17, 2018.

Mir17. Ilya Mironov. Rényi differential privacy. In *IEEE Computer Security Foundations Symposium*, 2017.

MKMS22. Jose Maria Bermudo Mera, Angshuman Karmakar, Tilen Marc, and Azam Soleimanian. Efficient lattice-based inner-product functional encryption. In *PKC 2022, Part II*, LNCS, pages 163–193. Springer, Heidelberg, May 2022.

MP13. Daniele Micciancio and Chris Peikert. Hardness of SIS and LWE with small parameters. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 21–39. Springer, Heidelberg, August 2013.

MR04. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 372–381, 2004.

Pei09. Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 333–342. ACM Press, May / June 2009.

PS24. Alain Passelègue and Damien Stehlé. Low communication threshold fully homomorphic encryption, 2024. To appear in ASIACRYPT.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

Reg09. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, 2009.

RSW18. Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the ring-LWE and polynomial-LWE problems. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 146–173. Springer, Heidelberg, April / May 2018.

STA20. Chao Sun, Mehdi Tibouchi, and Masayuki Abe. Revisiting the hardness of binary error LWE. In Joseph K. Liu and Hui Cui, editors, *ACISP 20*, volume 12248 of *LNCS*, pages 425–444. Springer, Heidelberg, November / December 2020.

vEH14. Tim van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.