Adaptive Leaky LWE and Its Security

Chuanqi Zhang, Muhammed F. Esgin and Ron Steinfeld

Monash University, Melbourne, Victoria, Australia

Abstract. The Leaky LWE problem, proposed by Lai, Swarnakar, and Woo (CiC '25), studies whether LWE samples remain pseudorandom in the presence of noisy linear information about their secrets and errors. In Leaky LWE, the leakage matrix may depend on the public matrix, but must be chosen *before* the LWE challenge is revealed. We study t-round *adaptive* Leaky LWE, where each leakage matrix may depend on the challenge and all preceding leakage answers. Under explicit norm, smoothing, and sampling conditions, we reduce Gaussian-secret LWE to this problem at the corresponding source parameters for polynomially many rounds. The reduction preserves the public-matrix distribution and access to its correlated oracle, and supports rectangular leakage matrices and non-spherical discrete Gaussians over arbitrary orders in number fields.

Keywords: LWE · Hint-MLWE · Leaky LWE · Adaptive Leaky LWE

1 Introduction

The Learning with Errors (LWE) problem [Reg09] has been a cornerstone of lattice-based cryptography. It asks to distinguish samples of the form

$$\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e} \pmod{q}$$

from uniformly random $\mathbf{b}' \xleftarrow{\$} U(\mathbb{Z}_q^m)$, given the matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$. In the standard formulation, $\mathbf{A}$ is uniformly random, the secret $\mathbf{s}$ is sampled uniformly or from a suitable discrete Gaussian distribution, and the error $\mathbf{e}$ is sampled from a narrow discrete Gaussian distribution. The problem extends naturally to algebraic lattices. For an order R in a number field and $R_q := R/qR$, taking the secret to be a single element of R gives Ring-LWE [SSTX09, LPR13], and taking the secret to be a vector in R^n and the public matrix over R_q gives Module-LWE [BGV14, LS15]. Thus, the formulation $\mathbf{A} \in R_q^{n\times m}$, $\mathbf{s} \in R^n$, and $\mathbf{e} \in R^m$ used in this paper includes plain LWE, Ring-LWE, and Module-LWE as special cases.

Noisy linear information in LWE-based proofs. A growing line of work studies whether an LWE sample remains pseudorandom when the distinguisher is additionally given noisy linear information about its secret and/or error [DKL+23, KLSS23, ENP24, LSW25, WSS+26, ESTX26, HKM+26]. We use column vectors throughout. Such leakage information has the form

$$\boldsymbol{\ell} = \mathbf{L}^\top \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} + \mathbf{f}, \tag{1}$$

where $\mathbf{L} \in R^{(n+m)\times Q}$ is a low-norm matrix and $\mathbf{f} \in R^Q$ is an independent masking error. The LWE equation is evaluated modulo q, whereas the leakage is evaluated over R.

E-mail: chuanqi.zhang@monash.edu (Chuanqi Zhang), muhammed.esgin@monash.edu (Muhammed F. Esgin), ron.steinfeld@monash.edu (Ron Steinfeld)

Relations of this form arise, for example, from responses in lattice-based proof systems and from partial decryptions or ciphertext components in threshold and laconic cryptography. A classical way to hide their dependence on $(\mathbf{s}, \mathbf{e})$ is statistical noise flooding. This requires the masking error to dominate the leaked term by a super-polynomial factor and consequently forces a super-polynomial modulus in typical applications. Error-Leakage LWE, Hint-MLWE, and Leaky LWE provide a computational alternative to noise flooding, proving that the LWE sample remains pseudorandom even given this noisy linear information [DKL+23, KLSS23, LSW25].

The adaptivity gap. Lai, Swarnakar, and Woo formalized this approach through Leaky LWE [LSW25]. Their experiment permits the public data to be sampled as $(\mathbf{A}, \mathcal{O}_{\mathbf{A}}) \xleftarrow{\$} \mathcal{D}_{\mathbf{A}}$, where $\mathcal{O}_{\mathbf{A}}$ is an oracle correlated with $\mathbf{A}$[1], and permits the distinguisher, after seeing $\mathbf{A}$ and with access to $\mathcal{O}_{\mathbf{A}}$, to choose an admissible leakage matrix $\mathbf{L}$. The matrix $\mathbf{L}$, however, must be selected *before* the LWE challenge $\mathbf{b}$ is generated. In some security proofs the leakage coefficients become known only after the public LWE sample has been fixed. They may be derived from an adversarially selected ciphertext, protocol message, or query depending on the complete public transcript, and later coefficients may further depend on earlier leakage answers. These possibilities give the following timings:

$$\begin{aligned} \text{(one-shot) semi-adaptive:} &\quad \mathbf{A} \to \mathbf{L} \to (\mathbf{b}, \boldsymbol{\ell}), \\ \text{(one-shot) adaptive:} &\quad (\mathbf{A}, \mathbf{b}) \to \mathbf{L} \to \boldsymbol{\ell}, \\ \text{(t-round) adaptive:} &\quad (\mathbf{A}, \mathbf{b}) \to \mathbf{L}_1 \to \boldsymbol{\ell}_1 \to \cdots \to \mathbf{L}_t \to \boldsymbol{\ell}_t. \end{aligned}$$

In the one-shot adaptive game, $\mathbf{L}$ may depend on the challenge $\mathbf{b}$, which is itself correlated with the leaked secret and error in the real branch. In the t-round adaptive game, $\mathbf{L}_i$ may additionally depend on all preceding leakage answers. Consequently, the semi-adaptive version cannot provide security for either of the latter two timings. This motivates the question addressed in this work:

Does the corresponding LWE problem also imply the hardness of adaptive Leaky LWE while preserving the general leakage and public-matrix setting of Leaky LWE?

1.1 Our Contribution

We answer this question affirmatively. We define t-round adaptive Leaky LWE, denoted $\mathsf{ALLWE}[t]$, with the one-shot game as the case $t = 1$. Our main theorem turns every adaptive distinguisher into a Leaky LWE distinguisher. Combining it with [LSW25, Theorem 3] gives the reduction from Gaussian-secret LWE.

More precisely, under the parameter conditions stated in Theorem 1 and Corollary 1, our reduction converts every admissible PPT adversary against t-round adaptive Leaky LWE into a PPT distinguisher for Gaussian-secret LWE at the corresponding source parameters. The advantage loss consists of a one-time loss from the Leaky LWE reduction and a per-round simulation error that accumulates linearly in t. When these errors are negligible, the total loss remains negligible for polynomially many leakage rounds. The one-shot adaptive result follows by setting $t = 1$.

The reduction preserves the public distribution. For any specified distribution $(\mathbf{A}, \mathcal{O}_{\mathbf{A}}) \xleftarrow{\$} \mathcal{D}_{\mathbf{A}}$ for which the source LWE problem is hard, the source and target experiments use the same $\mathcal{D}_{\mathbf{A}}$ and the same live correlated oracle $\mathcal{O}_{\mathbf{A}}$. Note that this is a distribution-preserving statement, rather than asserting that LWE is hard for every possible choice of $\mathcal{D}_{\mathbf{A}}$.

[1]The oracle may depend on auxiliary information generated together with $\mathbf{A}$, such as a trapdoor.

Concurrent work. Independently and concurrently with this work, Hwang, Katsumata, Min, Niot, and Song study post-challenge adaptive hints under the name Adaptive Hint-MLWE [HKM+26]. Their distinguisher receives an MLWE challenge and then makes adaptive ring-vector queries, each returning one noisy ring-valued linear function of the same short vector. Their experiment agrees with our scalar-hint special case for uniform public matrices over a power-of-two cyclotomic ring, an empty correlated oracle, and spherical Gaussians. Our theorem covers the more general public distribution, leakage matrix, and Gaussian setting of Leaky LWE. Section 5 compares the parameter conditions of the two reductions in their common setting, where neither set of sufficient conditions uniformly improves on the other.

1.2 Technical Overview

Writing $\mathbf{x} := \binom{\mathbf{s}}{\mathbf{e}}$, the LWE sample and leakage are

$$\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e} \pmod q, \quad \boldsymbol{\ell} = \mathbf{L}^\top \mathbf{x} + \mathbf{f}.$$

The obstacle is that the adversary chooses $\mathbf{L}_i$ only after receiving $(\mathbf{A}, \mathbf{b})$, while a semi-adaptive adversary must submit its leakage matrix before $\mathbf{b}$ is generated. We remove this obstacle by asking the Leaky LWE challenger for independently noised copies of the entire vector $\mathbf{x}$, without anticipating the later matrices.

For t target queries, the reduction submits the single repeated-identity matrix

$$\mathbf{J}_t = [\, I_{n+m} \mid \cdots \mid I_{n+m} \,] \in R^{(n+m)\times t(n+m)}$$

to the Leaky LWE experiment. Its answer is a vector in $R^{t(n+m)}$, which we split into t blocks of length $n+m$. Denote the ith block by $\mathbf{z}_i = \mathbf{x} + \mathbf{f}_{\mathsf{LLWE},i} \in R^{n+m}$, where $\mathbf{f}_{\mathsf{LLWE},1}, \ldots, \mathbf{f}_{\mathsf{LLWE},t} \in R^{n+m}$ are the independent blocks of the semi-adaptive leakage noise $\mathbf{f}_{\mathsf{LLWE}}$. In round i, the adversary chooses $\mathbf{L}_i$ based on $(\mathbf{A}, \mathbf{b})$ and the preceding leakage answers. The reduction samples an additional correction noise $\mathbf{h}_i \in R^Q$ and returns

$$\widetilde{\boldsymbol{\ell}}_i = \mathbf{L}_i^\top \mathbf{z}_i + \mathbf{h}_i = \mathbf{L}_i^\top \mathbf{x} + (\mathbf{L}_i^\top \mathbf{f}_{\mathsf{LLWE},i} + \mathbf{h}_i).$$

It remains to choose the distribution of the correction noise $\mathbf{h}_i$ so that the parenthesized term has the target leakage-error distribution.

In Euclidean coordinates, let $T_{\mathbf{L}_i}$ denote the real linear map representing $\mathbf{L}_i^\top$. Let $\sigma_{f_{\mathsf{LLWE}}}$ be the Gaussian width of each identity-query noise block, and let $\Sigma_{\mathbf{f}}$ be the Gaussian parameter matrix of the target adaptive leakage error $\mathbf{f}_i$. For the correction noise $\mathbf{h}_i$, we set

$$\Sigma_{\mathbf{h}}(\mathbf{L}_i) := \Sigma_{\mathbf{f}} - \sigma^2_{f_{\mathsf{LLWE}}} T_{\mathbf{L}_i} T^\top_{\mathbf{L}_i}.$$

The norm restriction on the admissible leakage family guarantees that this matrix is positive definite with a uniform spectral gap. Under the efficient-sampling condition in Theorem 1, a discrete Gaussian with this parameter can be sampled after $\mathbf{L}_i$ is known using Peikert's generic sampler [Pei10]. Furthermore, applying [GMPW20, Theorem 1] shows, under our smoothing condition, that $\mathbf{L}_i^\top \mathbf{f}_{\mathsf{LLWE},i} + \mathbf{h}_i$ is statistically close to the target discrete Gaussian with parameter $\Sigma_{\mathbf{f}}$.

Each round uses a different noisy copy $\mathbf{z}_i = \mathbf{x} + \mathbf{f}_{\mathsf{LLWE},i}$. Reusing the same copy would make the simulated leakage errors share the same underlying noise, but the real game uses independent leakage errors in every round. Since $\mathbf{f}_{\mathsf{LLWE},1}, \ldots, \mathbf{f}_{\mathsf{LLWE},t}$ are independent and $\mathbf{f}_{\mathsf{LLWE},i}$ has not been used in any preceding answer, it is independent of both the preceding transcript and the chosen matrix $\mathbf{L}_i$. The single-round analysis therefore applies to each adaptive query. A hybrid argument over the t answers gives the advantage loss in Theorem 1.

Finally, Corollary 1 shows that the repeated-identity query satisfies the norm bound required by [LSW25, Condition 1]. Applying their theorem once then gives the reduction from LWE.

1.3 Overview of LWE with Hints

Noisy hints on the secret and error. Döttling et al. introduced Error-Leakage LWE [DKL+23]. In addition to the LWE sample $\mathbf{b} = \mathbf{A}^\top \mathbf{s} + \mathbf{e} \pmod q$, the distinguisher is given a noisy linear function of the error,

$$\boldsymbol{\ell}_e = \mathbf{L}_e^\top \mathbf{e} + \mathbf{f},$$

where $\mathbf{L}_e \in R^{m \times Q}$ is an admissible low-norm matrix chosen after $\mathbf{A}$ is revealed and $\mathbf{f} \in R^Q$ is an independent masking error. Thus, Error-Leakage LWE permits leakage on $\mathbf{e} \in R^m$, but not directly on the uniform secret $\mathbf{s} \in R_q^n$.

Kim et al. introduced Hint-MLWE [KLSS23], which also permits noisy information about the secret. Its hints, however, have a particular structured form. The challenger samples a tuple of public ring elements $(\gamma_1, \ldots, \gamma_Q) \in R^Q$ according to a prescribed joint distribution and returns

$$\boldsymbol{\ell} = \begin{bmatrix} \gamma_1 I_{n+m} \\ \vdots \\ \gamma_Q I_{n+m} \end{bmatrix} \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} + \mathbf{f},$$

where $\mathbf{f} \in R^{(n+m)Q}$ stacks Q independently sampled masking-error vectors. Thus each γ_i acts by ring multiplication on every coordinate of the joint secret-error vector, and the tuple is sampled by the challenger rather than selected by the distinguisher.

Espitau, Niot, and Prest generalized this formulation to MatrixHint-MLWE [ENP24], allowing general square matrices in the hints. They write the secret $\mathbf{s} \in \mathbb{Z}^{\varphi n}$ and error $\mathbf{e} \in \mathbb{Z}^{\varphi m}$ as vectors of their integer coefficients, where φ is the degree of the number field. The challenger samples public low-norm matrices $\mathbf{L}_1, \ldots, \mathbf{L}_Q \in \mathbb{Z}^{\varphi(n+m) \times \varphi(n+m)}$ according to a prescribed joint distribution and returns

$$\boldsymbol{\ell} = \begin{bmatrix} \mathbf{L}_1^\top \\ \vdots \\ \mathbf{L}_Q^\top \end{bmatrix} \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} + \mathbf{f},$$

where $\mathbf{f} \in \mathbb{Z}^{\varphi(n+m)Q}$ stacks Q independently sampled masking-error vectors.

Wang et al. subsequently considered rectangular hint matrices over R [WSS+26]. The challenger samples a public low-norm matrix $\mathbf{L} \in R^{(n+m) \times Q}$ from a prescribed distribution and returns

$$\boldsymbol{\ell} = \mathbf{L}^\top \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} + \mathbf{f},$$

where $\mathbf{s} \in R^n$ and $\mathbf{e} \in R^m$ are the secret and error vectors, and $\mathbf{f} \in R^Q$ is an independent masking error.

Lai, Swarnakar, and Woo introduced Leaky LWE [LSW25] as a common generalization of Error-Leakage LWE and Hint-MLWE. Its leakage also has the general form as in (1), where $\mathbf{L} \in R^{(n+m) \times Q}$ is an admissible leakage matrix and $\mathbf{f} \in R^Q$ is an independent leakage error. Their experiment follows this template:

1. The challenger samples public data $(\mathbf{A}, \mathcal{O}_\mathbf{A}) \xleftarrow{\$} \mathcal{D}_\mathbf{A}$, where $\mathcal{O}_\mathbf{A}$ may be correlated with $\mathbf{A}$. For example, $\mathbf{A}$ may be the output of a random oracle at a fixed public input, and $\mathcal{O}_\mathbf{A}$ may provide continued access to that same oracle. The source LWE assumption is then understood with access to this oracle.

2. After seeing $\mathbf{A}$ and with access to $\mathcal{O}_\mathbf{A}$, but before the LWE challenge is generated, the adversary chooses the leakage matrix $\mathbf{L}$.

3. The challenger samples the LWE secret $\mathbf{s} \xleftarrow{\$} \mathcal{D}_{\mathbf{s}}$, the LWE error $\mathbf{e} \xleftarrow{\$} \mathcal{D}_{\mathbf{e}}$, and an independent leakage error $\mathbf{f} \xleftarrow{\$} \mathcal{D}_{\mathbf{f}}$. It lets $\mathbf{b}$ be either $\mathbf{A}^\top \mathbf{s} + \mathbf{e} \pmod q$ or a uniformly random element of R_q^m, and forms the leakage as in (1).
4. Given $(\mathbf{A}, \mathbf{b}, \boldsymbol{\ell})$ and continued access to the same oracle $\mathcal{O}_{\mathbf{A}}$, the adversary distinguishes the two choices of $\mathbf{b}$. In both branches the leakage is computed from the same secret, error, and leakage-error distributions.

Their theorem proves Leaky LWE hard under the corresponding LWE assumption, preserving the same specified $\mathcal{D}_{\mathbf{A}}$ and correlated oracle under explicit Gaussian conditions. It also revisits the Error-Leakage LWE argument and supplies corrected conditions. The experiment is nevertheless semi-adaptive because $\mathbf{L}$ is fixed before $\mathbf{b}$ is generated.

Interactive variants. Esgin et al. introduced Hint-MLWE-PRF as an intermediate problem in their study of MLWE-PRF with re-use [ESTX26]. Focusing on its MLWE-hint component, a single secret $\mathbf{s} \in R^n$ is fixed throughout the experiment. For each query i, the adversary first supplies an admissible low-norm matrix $\mathbf{L}_i \in R^{m \times Q}$; the challenger then samples a fresh matrix $\mathbf{A}_i \in R_q^{n \times m}$ and a fresh error $\mathbf{e}_i \in R^m$, and returns

$$\mathbf{b}_i = \mathbf{A}_i^\top \mathbf{s} + \mathbf{e}_i \pmod q$$

together with a noisy error hint, which in our notation has the form

$$\boldsymbol{\ell}_i = \mathbf{L}_i^\top \mathbf{e}_i + \mathbf{f}_i,$$

where $\mathbf{f}_i \in R^Q$ is a fresh masking error. Thus, the samples $\mathbf{b}_i$ provide repeated noisy information about the same secret $\mathbf{s}$, while the error $\mathbf{e}_i$ changes from one query to the next. The queries are interactive: $\mathbf{L}_i$ may depend on the preceding transcript, but it must be chosen before the corresponding fresh sample $(\mathbf{A}_i, \mathbf{b}_i)$ is generated. Their full Hint-MLWE-PRF experiment additionally supplies linked noisy hints on PRF errors, and it is therefore distinct from adaptive leakage on one fixed joint vector $(\mathbf{s}, \mathbf{e})$.

Knapsack and coset-error variants. Kume et al. formalized the natural knapsack variant of MatrixHint-MLWE [KSSY26]. Writing its short knapsack vector in the same joint form $\binom{\mathbf{s}}{\mathbf{e}}$, the challenger samples independent uniform matrices $\mathbf{A}_s \in R_q^{n \times m}$ and $\mathbf{A}_e \in R_q^{m \times m}$ and a public hint matrix $\mathbf{L} \in R_q^{(n+m) \times Q}$, and returns

$$\mathbf{b} = \mathbf{A}_s^\top \mathbf{s} + \mathbf{A}_e^\top \mathbf{e} \pmod q,$$

together with the matrix hint

$$\boldsymbol{\ell} = \mathbf{L}^\top \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} + \mathbf{f},$$

where $\mathbf{f} \in R^Q$ is an independent short masking error and $\boldsymbol{\ell} \in R^Q$. In plain LWE, the error block is fixed as $\mathbf{A}_e = I_m$, while in the knapsack formulation both $\mathbf{A}_s$ and $\mathbf{A}_e$ are uniformly random. The Knapsack MatrixHint-MLWE assumption states that $\mathbf{b}$ remains pseudorandom even given $(\mathbf{A}_s, \mathbf{A}_e, \mathbf{L}, \boldsymbol{\ell})$.

Kume et al. further introduced Coset Error Knapsack MatrixHint-MLWE [KSSY26], which changes the sampling of the short knapsack vector. First, the challenger samples and publishes $(\mathbf{A}_s, \mathbf{A}_e, \mathbf{L})$. After seeing these matrices, the adversary chooses $\mathbf{c} \in R_q^{n+m}$, thereby specifying a coset of a prescribed q-ary lattice $\Lambda \subseteq R^{n+m}$. Because Λ is q-ary, any lift of $\mathbf{c}$ to R^{n+m} defines the same coset, which we write as $\Lambda + \mathbf{c}$. The challenger then samples

$$\begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} \xleftarrow{\$} D_{\Lambda + \mathbf{c}, \sigma I_{\varphi(n+m)}}.$$

It also samples an independent short masking error $\mathbf{f} \in R^Q$ and returns

$$\mathbf{b} = \mathbf{A}_s^\top \mathbf{s} + \mathbf{A}_e^\top \mathbf{e} \pmod q,$$

together with

$$\boldsymbol{\ell} = \mathbf{L}^\top \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} + \mathbf{f}.$$

As before, the random experiment replaces $\mathbf{b}$ with a uniformly random vector while leaving $\boldsymbol{\ell}$ unchanged. Thus, relative to Knapsack MatrixHint-MLWE, the new feature is the adversarially selected lattice coset from which the short vector is sampled.

Exact and more general leakage. A related line studies exact rather than noisy linear leakage. Extended LWE [ASP12], multi-hint Extended LWE [ALS16], and Ring/Module-LWE under linear leakage [WLL24] admit reductions from appropriate standard LWE problems, but restrict the amount or form of exact leakage. In the latter formulation, the exact leakage functions are committed before the public matrix and MLWE sample are drawn.

Besides, Boudgoust et al. study search Module-LWE with general secret-error distributions and leakage functions by analyzing the entropy remaining after leakage [BJTW26]. Their results concern recovery of the secret, whereas we prove that the LWE sample remains pseudorandom in the presence of adaptively chosen noisy linear hints.

Under their respective algebraic settings, parameter conditions, and security models, the above works give reductions from appropriate LWE problems to the corresponding hinted problems. These assumptions have supported laconic cryptography [DKL+23], practical lattice-based proofs [KLSS23], threshold signatures [ENP24, DPKM+24, BKL+25], multi-message public-key encryption [WSS+26], post-quantum oblivious pseudorandom functions [ESTX26], and succinct zero-knowledge arguments [KSSY26].

Organization. Section 2 introduces the notation, summarized in Table 1, and defines the security experiments. Section 3 establishes the hardness of adaptive Leaky LWE via reductions from Leaky LWE and therefore from LWE. Section 4 gives explicit parameter bounds and specializes the leakage matrix to several common hint forms. Section 5 compares our reduction with that of Hwang et al. [HKM+26] under the same query count, query-norm bound, and advantage-loss budget, expressing the leakage-noise width relative to the source-LWE width.

2 Preliminaries and Security Games

All objects below are ensembles indexed by a security parameter λ. Dimensions, moduli, round counts, algorithms, and encodings are polynomially bounded unless stated otherwise.

2.1 Ring coordinates, Euclidean coordinates, and Gaussians

Let $\mathbb{Q}$ denote the field of rational numbers, and let K be a number field of degree φ over $\mathbb{Q}$. Let $R \subseteq K$ be an order, and put $R_q := R/qR$. Denote by $U(R_q)$ the uniform distribution on R_q. Fix a $\mathbb{Q}$-linear embedding

$$\psi_1 : K \to \mathbb{R}^\varphi$$

that maps an integral basis to φ linearly independent real vectors. The coefficient and realified canonical embeddings are standard examples. Extend ψ_1 coordinatewise:

$$\psi_k(x_1, \ldots, x_k) := \big(\psi_1(x_1), \ldots, \psi_1(x_k)\big) \in \mathbb{R}^{\varphi k},$$

Table 1: Summary of notation. A superscript $*$ on LWE parameters usually denotes the source instance. Gaussian scales are widths in the convention $\exp(-\pi\|\mathbf{x}\|_2^2/\sigma^2)$.

Notation	Description	Reference
Setting and security games		
λ	Security parameter.	§ 2
$R,\ R_q,\ q$	Number-field order, quotient $R_q = R/qR$, and modulus.	§ 2.1
$U(R_q)$	Uniform distribution on R_q.	§ 2.1
$\varphi,\ N$	Number-field degree; $N = \varphi$ in the power-of-two cyclotomic specializations.	§ 2.1; § 4.1
$n,\ m$	Secret and error dimensions; $\mathbf{A} \in R_q^{n\times m}$.	§ 2.2
$Q,\ t$	Ring elements in each leakage answer, and number of adaptive rounds, respectively.	§ 2.2; § 2.4
$\mathcal{D}_{\mathbf{A}},\ \mathcal{O}_{\mathbf{A}}$	Joint distribution of the public matrix and its correlated oracle; the oracle itself.	§ 2.2
$\mathbf{s},\ \mathbf{e},\ \mathbf{x}$	Secret in R^n, error in R^m, and $\mathbf{x} = (\mathbf{s}^\top, \mathbf{e}^\top)^\top$.	§ 2.2
$\mathbf{L},\ \mathbf{L}_i,\ \mathcal{L}$	Leakage matrix in $R^{(n+m)\times Q}$, query in round i, and admissible query family.	§ 2.2; § 2.4
$\mathsf{p}^*,\ \mathsf{p}$	Source-LWE and target-Leaky LWE parameter sets.	§ 2.2
$\mathsf{LLWE},\ \mathsf{ALLWE}[t]$	Leaky LWE and t-round adaptive Leaky LWE experiments.	§ 2.3; § 2.4
Lattices and Gaussian parameters		
$\psi_k,\ \Lambda_k$	Euclidean embedding of K^k and the lattice $\Lambda_k = \psi_k(R^k)$.	§ 2.1
$\Psi,\ T_{\mathbf{L}}$	Real matrix representation of ring matrices; $T_{\mathbf{L}} = \Psi(\mathbf{L}^\top)$.	§ 2.1
$I_k,\ \mathbf{J}_t$	$k \times k$ identity matrix over the ambient ring or field; horizontal concatenation of t copies of I_{n+m}.	§ 2.1; § 3.1
$D_{R^k,\sqrt{\Sigma}},\ \Sigma$	Discrete Gaussian on R^k and its Gaussian parameter matrix; spherical width σ means $\Sigma = \sigma^2 I_{\varphi k}$.	§ 2.1
$s_{\min}(\Sigma),\ s_{\max}(\Sigma)$	Smallest and largest eigenvalues for $\Sigma \succ 0$, respectively.	§ 2.1
$\eta_\varepsilon(\Lambda)$	Smoothing parameter of lattice Λ at tolerance ε.	§ 2.1
$\mathbf{B}_Q$	Public basis matrix with $\Lambda_Q = \mathbf{B}_Q\mathbb{Z}^{\varphi Q}$.	§ 3.1
Noise and reduction parameters		
$\mathbf{f}_i,\ \Sigma_{\mathbf{f}},\ \sigma_f$	Fresh target leakage noise in R^Q with distribution $\mathcal{D}_{\mathbf{f}}$. In Theorem 1, its Gaussian parameter is $\Sigma_{\mathbf{f}}$; for spherical noise, $\Sigma_{\mathbf{f}} = \sigma_f^2 I_{\varphi Q}$.	§ 3.1; § 4.1
$\mathbf{f}_{\mathsf{LLWE}},\ \sigma_{f_{\mathsf{LLWE}}}$	Auxiliary Leaky LWE noise in $R^{t(n+m)}$ and its spherical width. Its ith block $\mathbf{f}_{\mathsf{LLWE},i}$ lies in R^{n+m}.	(3); (6)
$\mathbf{h},\ \Sigma_{\mathbf{h}}(\mathbf{L})$	Ideal correction noise in R^Q and its Gaussian parameter, chosen for query $\mathbf{L}$.	(5)
$\varepsilon_0,\ \varepsilon_c,\ \varepsilon_s$	Tolerances for the reduction from LWE to Leaky LWE, Gaussian correction, and sampler smoothing, respectively.	§ 3.1; § 3.3
ν	Finite-precision and truncation error of the correction-noise sampler.	Lemma 2
r_Q	Gaussian rounding width from the correction-noise sampler.	§ 3.1
κ	Positive eigenvalue lower bound controlling the precision of the correction-noise sampler.	§ 3.1
Spherical specializations and comparison		
$\sigma_s^*,\ \sigma_e^*,\ \sigma_s,\ \sigma_e$	Source and target secret/error widths, respectively.	§ 4.1
$\sigma^*,\ \sigma$	Common secret/error widths of the source and target in Section 5. The comparison uses σ_f/σ^*.	§ 5
B	Upper bound on $\|T_{\mathbf{L}}\|_2^2$.	§ 4.1; § 5

and write

$$\Lambda_k := \psi_k(R^k) \subseteq \mathbb{R}^{\varphi k},$$

which is a full-rank lattice. There is a unique multiplication representation $\Psi : K \to \mathbb{R}^{\varphi\times\varphi}$ satisfying

$$\psi_1(ax) = \Psi(a)\psi_1(x) \text{ for } a, x \in K.$$

We extend Ψ blockwise to ring matrices. Thus, for $\mathbf{L} \in R^{a\times b}$ and $\mathbf{x} \in K^b$,

$$\psi_a(\mathbf{L}\mathbf{x}) = \Psi(\mathbf{L})\psi_b(\mathbf{x}), \quad \Psi(\mathbf{L}) \in \mathbb{R}^{\varphi a\times\varphi b}.$$

We also write

$$T_{\mathbf{L}} := \Psi(\mathbf{L}^\top) \in \mathbb{R}^{\varphi b\times\varphi a},$$

for the real matrix representing the map $\mathbf{x} \mapsto \mathbf{L}^\top\mathbf{x}$. For a full-rank lattice $\Lambda \subseteq \mathbb{R}^N$, a positive-definite matrix $\Sigma \in \mathbb{R}^{N\times N}$, and a center $\mathbf{c} \in \mathbb{R}^N$, write

$$\rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbf{y}) := \exp\bigl(-\pi(\mathbf{y}-\mathbf{c})^\top\Sigma^{-1}(\mathbf{y}-\mathbf{c})\bigr)$$

and, for a countable set S, put $\rho_{\sqrt{\Sigma},\mathbf{c}}(S) := \sum_{\mathbf{y}\in S}\rho_{\sqrt{\Sigma},\mathbf{c}}(\mathbf{y})$. Let $D_{\Lambda+\mathbf{c},\sqrt{\Sigma}}$ denote the zero-centered discrete Gaussian on the coset $\Lambda+\mathbf{c}$, defined by

$$D_{\Lambda+\mathbf{c},\sqrt{\Sigma}}(\mathbf{y}) := \frac{\rho_{\sqrt{\Sigma},\mathbf{0}}(\mathbf{y})}{\rho_{\sqrt{\Sigma},\mathbf{0}}(\Lambda+\mathbf{c})} \text{ for } \mathbf{y} \in \Lambda+\mathbf{c}.$$

When $\mathbf{c} = 0$, we write $D_{\Lambda,\sqrt{\Sigma}}$, and use $D_{R^k,\sqrt{\Sigma}}$ for the inverse image of $D_{\Lambda_k,\sqrt{\Sigma}}$ under ψ_k. For $\sigma > 0$, taking $\Sigma = \sigma^2 I_{\varphi k}$ gives the spherical discrete Gaussian $D_{R^k,\sigma I_{\varphi k}}$ of width σ. Write

$$\Lambda^* := \{\mathbf{y} \in \operatorname{span}(\Lambda) : \langle\mathbf{y},\mathbf{x}\rangle \in \mathbb{Z} \text{ for every } \mathbf{x} \in \Lambda\}$$

for the dual lattice of Λ. For $\varepsilon > 0$, the smoothing parameter is

$$\eta_\varepsilon(\Lambda) := \inf\left\{ s > 0 : \sum_{\mathbf{y}\in\Lambda^*\setminus\{\mathbf{0}\}} \exp(-\pi s^2\|\mathbf{y}\|_2^2) \le \varepsilon \right\}.$$

For real symmetric matrices A and B of the same size, we write $A \preceq B$ (or equivalently, $B \succeq A$) if $B - A$ is positive semidefinite, and $A \succ 0$ if A is positive definite. For a positive-definite matrix Σ, let $\sqrt{\Sigma}$ denote the unique symmetric positive-definite square root of Σ, and let $s_{\min}(\Sigma)$ and $s_{\max}(\Sigma)$ denote its smallest and largest eigenvalues, respectively. We write I_k for the $k \times k$ identity matrix over the ambient ring or field.

2.2 The public-matrix distribution and its oracle

Let $\mathcal{D}_{\mathbf{A}}$ be a public distribution over pairs

$$(\mathbf{A}, \mathcal{O}_{\mathbf{A}}) \xleftarrow{\$} \mathcal{D}_{\mathbf{A}}, \quad \mathbf{A} \in R_q^{n\times m}.$$

The adversary has access throughout the security experiment to the same oracle $\mathcal{O}_{\mathbf{A}}$, which may be correlated with $\mathbf{A}$. All secret, error, and leakage-noise samples below are sampled independently of $(\mathbf{A}, \mathcal{O}_{\mathbf{A}})$. Uniform $\mathbf{A}$ and an empty oracle form only one special case.

The source-LWE parameter set is

$$\mathsf{p}^* := (R, n, m, q, \mathcal{D}_{\mathbf{A}}, \mathcal{D}^*_{\mathbf{s}}, \mathcal{D}^*_{\mathbf{e}}),$$

where $\mathcal{D}^*_{\mathbf{s}}$ is over R^n and $\mathcal{D}^*_{\mathbf{e}}$ is over R^m. A target Leaky LWE parameter set is

$$\mathsf{p} := (R, n, m, q, Q, \mathcal{D}_{\mathbf{A}}, \mathcal{D}_{\mathbf{s}}, \mathcal{D}_{\mathbf{e}}, \mathcal{D}_{\mathbf{f}}, \mathcal{L}),$$

where $\mathcal{D}_{\mathbf{s}}, \mathcal{D}_{\mathbf{e}}, \mathcal{D}_{\mathbf{f}}$ are distributions over R^n, R^m, R^Q, respectively, and

$$\mathcal{L} \subseteq R^{n\times Q} \times R^{m\times Q}$$

is polynomial-time decidable. For $(\mathbf{L}_s, \mathbf{L}_e) \in \mathcal{L}$, put

$$\mathbf{x} := \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} \in R^{n+m}, \quad \mathbf{L} := \begin{pmatrix} \mathbf{L}_s \\ \mathbf{L}_e \end{pmatrix} \in R^{(n+m)\times Q}.$$

We identify each query pair $(\mathbf{L}_s, \mathbf{L}_e)$ with its stacked matrix $\mathbf{L}$, and use $\mathcal{L}$ for the admissible family in either representation. Then $\mathbf{L}^\top \mathbf{x} = \mathbf{L}_s^\top \mathbf{s} + \mathbf{L}_e^\top \mathbf{e} \in R^Q$. This leakage equality is evaluated over R, without reduction modulo q. Reduction modulo q is used only in the LWE equation.

2.3 Leaky LWE

Let $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$. Define $\mathrm{Exp}^{\mathsf{LLWE}}_{\mathsf{p},\mathcal{A},\beta}(1^\lambda)$ as follows.

1. Sample $(\mathbf{A}, \mathcal{O}_{\mathbf{A}}) \xleftarrow{\$} \mathcal{D}_{\mathbf{A}}$.
2. Before $\mathbf{b}$ is generated, run

   $$(\mathbf{L}_s, \mathbf{L}_e, \mathsf{st}) \xleftarrow{\$} \mathcal{A}_0^{\mathcal{O}_{\mathbf{A}}}(\mathbf{A}).$$

   Return zero if $(\mathbf{L}_s, \mathbf{L}_e) \notin \mathcal{L}$.
3. Independently sample

   $$\mathbf{s} \xleftarrow{\$} \mathcal{D}_{\mathbf{s}}, \quad \mathbf{e} \xleftarrow{\$} \mathcal{D}_{\mathbf{e}}, \quad \mathbf{f} \xleftarrow{\$} \mathcal{D}_{\mathbf{f}}.$$

   If $\beta = 0$, set $\mathbf{b} := \mathbf{A}^\top \mathbf{s} + \mathbf{e} \pmod q$; if $\beta = 1$, sample $\mathbf{b} \xleftarrow{\$} U(R_q^m)$.
4. Set

   $$\boldsymbol{\ell} := \mathbf{L}_s^\top \mathbf{s} + \mathbf{L}_e^\top \mathbf{e} + \mathbf{f} = \mathbf{L}^\top \mathbf{x} + \mathbf{f} \in R^Q.$$

5. Return $\mathcal{A}_1^{\mathcal{O}_{\mathbf{A}}}(\mathbf{b}, \boldsymbol{\ell}, \mathsf{st})$, with access to the same oracle instance and state as in the first stage.

The adversary's advantage is defined as

$$\mathrm{Adv}^{\mathsf{LLWE}}_{\mathsf{p},\mathcal{A}} := \left| \Pr[\mathrm{Exp}^{\mathsf{LLWE}}_{\mathsf{p},\mathcal{A},0} = 1] - \Pr[\mathrm{Exp}^{\mathsf{LLWE}}_{\mathsf{p},\mathcal{A},1} = 1] \right|.$$

For the experiment without leakage, we analogously write $\mathrm{Adv}^{\mathsf{LWE}}_{\mathsf{p}^*,\mathcal{A}}$ (and $\mathrm{Adv}^{\mathsf{MLWE}}_{\mathsf{p}^*,\mathcal{A}}$) for the distinguishing advantage of adversary $\mathcal{A}$ under the source parameters p^*.

2.4 Adaptive Leaky LWE

The adaptive game moves the choice of $\mathbf{L}$ past the LWE sample. In $\mathrm{Exp}^{\mathsf{ALLWE}[1]}_{\mathsf{p},\mathcal{A},\beta}$, sample $(\mathbf{A}, \mathcal{O}_{\mathbf{A}}), \mathbf{s}, \mathbf{e}$ exactly as above, form the real or uniform $\mathbf{b}$, and then run

$$(\mathbf{L}_s, \mathbf{L}_e, \mathsf{st}) \xleftarrow{\$} \mathcal{A}_0^{\mathcal{O}_{\mathbf{A}}}(\mathbf{A}, \mathbf{b}).$$

After checking admissibility, independently sample $\mathbf{f} \xleftarrow{\$} \mathcal{D}_{\mathbf{f}}$, set $\boldsymbol{\ell} = \mathbf{L}^\top \mathbf{x} + \mathbf{f}$, and return $\mathcal{A}_1^{\mathcal{O}_{\mathbf{A}}}(\boldsymbol{\ell}, \mathsf{st})$.

For a public round count $t = t(\lambda) \in \mathrm{poly}(\lambda)$, the t-round adaptive game $\mathrm{Exp}^{\mathsf{ALLWE}[t]}_{\mathsf{p},\mathcal{A},\beta}$ starts in the same way and gives $(\mathbf{A}, \mathbf{b})$ and access to $\mathcal{O}_{\mathbf{A}}$ to a stateful PPT adversary. For rounds $i = 1, \ldots, t$, it performs:

$\mathrm{Exp}^{\mathsf{LLWE}}_{\mathsf{p},\mathcal{A},\beta}(1^\lambda)$

$(\mathbf{A}, \mathcal{O}_{\mathbf{A}}) \xleftarrow{\$} \mathcal{D}_{\mathbf{A}}$

$(\mathbf{L}_s, \mathbf{L}_e, \mathsf{st}) \xleftarrow{\$} \mathcal{A}_0^{\mathcal{O}_{\mathbf{A}}}(\mathbf{A})$

if $(\mathbf{L}_s, \mathbf{L}_e) \notin \mathcal{L}$ **then return** 0

$\mathbf{s} \xleftarrow{\$} \mathcal{D}_{\mathbf{s}}; \quad \mathbf{e} \xleftarrow{\$} \mathcal{D}_{\mathbf{e}}; \quad \mathbf{f} \xleftarrow{\$} \mathcal{D}_{\mathbf{f}}$

if $\beta = 0$ **then** $\mathbf{b} \leftarrow \mathbf{A}^\top \mathbf{s} + \mathbf{e} \pmod q$

if $\beta = 1$ **then** $\mathbf{b} \xleftarrow{\$} U(R_q^m)$

$\mathbf{L} \leftarrow \begin{pmatrix} \mathbf{L}_s \\ \mathbf{L}_e \end{pmatrix}; \quad \mathbf{x} \leftarrow \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix}$

$\boldsymbol{\ell} \leftarrow \mathbf{L}^\top \mathbf{x} + \mathbf{f} \in R^Q$

$\beta' \leftarrow \mathcal{A}_1^{\mathcal{O}_{\mathbf{A}}}(\mathbf{b}, \boldsymbol{\ell}, \mathsf{st})$

return β'

Figure 1: The Leaky LWE security experiment. The leakage matrix is selected after $\mathbf{A}$ is revealed but before the LWE challenge $\mathbf{b}$ is generated.

$\mathrm{Exp}^{\mathsf{ALLWE}[1]}_{\mathsf{p},\mathcal{A},\beta}(1^\lambda)$

$(\mathbf{A}, \mathcal{O}_{\mathbf{A}}) \xleftarrow{\$} \mathcal{D}_{\mathbf{A}}$

$\mathbf{s} \xleftarrow{\$} \mathcal{D}_{\mathbf{s}}; \quad \mathbf{e} \xleftarrow{\$} \mathcal{D}_{\mathbf{e}}; \quad \mathbf{x} \leftarrow \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix}$

if $\beta = 0$ **then** $\mathbf{b} \leftarrow \mathbf{A}^\top \mathbf{s} + \mathbf{e} \pmod q$

if $\beta = 1$ **then** $\mathbf{b} \xleftarrow{\$} U(R_q^m)$

$(\mathbf{L}_s, \mathbf{L}_e, \mathsf{st}) \xleftarrow{\$} \mathcal{A}_0^{\mathcal{O}_{\mathbf{A}}}(\mathbf{A}, \mathbf{b})$

if $(\mathbf{L}_s, \mathbf{L}_e) \notin \mathcal{L}$ **then return** 0

$\mathbf{L} \leftarrow \begin{pmatrix} \mathbf{L}_s \\ \mathbf{L}_e \end{pmatrix}; \quad \mathbf{f} \xleftarrow{\$} \mathcal{D}_{\mathbf{f}}$

$\boldsymbol{\ell} \leftarrow \mathbf{L}^\top \mathbf{x} + \mathbf{f} \in R^Q$

$\beta' \leftarrow \mathcal{A}_1^{\mathcal{O}_{\mathbf{A}}}(\boldsymbol{\ell}, \mathsf{st})$

return β'

Figure 2: The adaptive Leaky LWE security experiment. The leakage matrix may depend on the complete LWE challenge $(\mathbf{A}, \mathbf{b})$.

(i) Based on its current state and the preceding answers, the adversary outputs $\mathbf{L}_i = \binom{\mathbf{L}_{s,i}}{\mathbf{L}_{e,i}} \in R^{(n+m)\times Q}$. Return zero if this matrix is not admissible.

(ii) Independently sample $\mathbf{f}_i \xleftarrow{\$} \mathcal{D}_{\mathbf{f}}$ and return

$$\boldsymbol{\ell}_i := \mathbf{L}_i^\top \mathbf{x} + \mathbf{f}_i \in R^Q.$$

After the last answer, the experiment returns the adversary's final bit. Thus the timing is

$$(\mathbf{A}, \mathbf{b}) \to \mathbf{L}_1 \to \boldsymbol{\ell}_1 \to \cdots \to \mathbf{L}_t \to \boldsymbol{\ell}_t.$$

The samples $\mathbf{f}_1, \ldots, \mathbf{f}_t$ are mutually independent and independent of all prior game randomness.

$\mathrm{Exp}^{\mathsf{ALLWE}[t]}_{\mathsf{p},\mathcal{A},\beta}(1^\lambda)$

1 : $(\mathbf{A}, \mathcal{O}_{\mathbf{A}}) \xleftarrow{\$} \mathcal{D}_{\mathbf{A}}$
2 : $\mathbf{s} \xleftarrow{\$} \mathcal{D}_{\mathbf{s}}$; $\mathbf{e} \xleftarrow{\$} \mathcal{D}_{\mathbf{e}}$; $\mathbf{x} \leftarrow \begin{pmatrix}\mathbf{s}\\ \mathbf{e}\end{pmatrix}$
3 : **if** $\beta = 0$ **then** $\mathbf{b} \leftarrow \mathbf{A}^\top \mathbf{s} + \mathbf{e} \pmod q$
4 : **if** $\beta = 1$ **then** $\mathbf{b} \xleftarrow{\$} U(R_q^m)$
5 : $i \leftarrow 0$; $\mathsf{bad} \leftarrow 0$
6 : $\beta' \leftarrow \mathcal{A}^{\mathcal{O}_{\mathbf{A}}, \mathsf{Leak}_{\mathbf{x}}}(\mathbf{A}, \mathbf{b})$
7 : **if** $\mathsf{bad} = 1 \ \vee \ i \neq t$ **then return** 0
8 : **return** β'

$\mathsf{Leak}_{\mathbf{x}}(\mathbf{L}_s, \mathbf{L}_e)$

1 : **if** $\mathsf{bad} = 1 \ \vee \ i = t$ **then**
  $\mathsf{bad} \leftarrow 1$; **return** $\perp$
2 : **if** $(\mathbf{L}_s, \mathbf{L}_e) \notin \mathcal{L}$ **then**
  $\mathsf{bad} \leftarrow 1$; **return** $\perp$
3 : $i \leftarrow i + 1$; $\mathbf{L} \leftarrow \begin{pmatrix}\mathbf{L}_s\\ \mathbf{L}_e\end{pmatrix}$
4 : $\mathbf{f}_i \xleftarrow{\$} \mathcal{D}_{\mathbf{f}}$
5 : $\boldsymbol{\ell}_i \leftarrow \mathbf{L}^\top \mathbf{x} + \mathbf{f}_i \in R^Q$
6 : **return** $\boldsymbol{\ell}_i$

Figure 3: The t-round adaptive Leaky LWE security experiment. Every query may depend on $(\mathbf{A}, \mathbf{b})$, the preceding leakage answers, and the transcript of the same live oracle $\mathcal{O}_{\mathbf{A}}$; every answer uses fresh independent leakage noise.

The adversary's advantage in the t-round adaptive experiment is defined as

$$\mathrm{Adv}^{\mathsf{ALLWE}[t]}_{\mathsf{p},\mathcal{A}} := \left| \Pr[\mathrm{Exp}^{\mathsf{ALLWE}[t]}_{\mathsf{p},\mathcal{A},0} = 1] - \Pr[\mathrm{Exp}^{\mathsf{ALLWE}[t]}_{\mathsf{p},\mathcal{A},1} = 1] \right|.$$

3 Hardness of Adaptive Leaky LWE

We simulate the t-round adaptive experiment using a semi-adaptive experiment with one fixed leakage query. The reduction obtains t independently noised copies of the secret-error vector and uses one copy to answer each adaptive query. We first prove this simulation, then combine it with [LSW25, Theorem 3].

3.1 Parameters and main theorem

Fix target parameters p as in Section 2.2, and let $t = t(\lambda)$ be polynomially bounded. We use the ring coordinates ψ_k, lattices Λ_k, and real matrix representation Ψ defined in Section 2.1. For an admissible leakage matrix $\mathbf{L} = \binom{\mathbf{L}_s}{\mathbf{L}_e}$, put $T_{\mathbf{L}} := \Psi(\mathbf{L}^\top)$, so that

$$T_{\mathbf{L}} \psi_{n+m}(\mathbf{x}) = \psi_Q(\mathbf{L}^\top \mathbf{x}) \text{ for every } \mathbf{x} \in R^{n+m}.$$

Thus $T_{\mathbf{L}} \in \mathbb{R}^{\varphi Q \times \varphi(n+m)}$ and $T_{\mathbf{L}} \Lambda_{n+m} \subseteq \Lambda_Q$. Let the target adaptive leakage-error distribution be

$$\mathcal{D}_{\mathbf{f}} = D_{R^Q, \sqrt{\Sigma_{\mathbf{f}}}}, \quad \Sigma_{\mathbf{f}} \in \mathbb{R}^{\varphi Q \times \varphi Q}, \quad \Sigma_{\mathbf{f}} \succ 0.$$

Let $\mathbf{B}_Q$ be the embedded matrix of a public $\mathbb{Z}$-basis of R^Q, so that $\Lambda_Q = \mathbf{B}_Q \mathbb{Z}^{\varphi Q}$.

Numerical representation. For the algorithms in this section and their specializations, public real scalars and matrix entries are uniformly computable to absolute precision 2^{-p} in time $\mathrm{poly}(\lambda, p)$. This requirement applies uniformly to $T_{\mathbf{L}}$ for admissible $\mathbf{L}$. We use known bounds of $2^{\mathrm{poly}(\lambda)}$ on their magnitudes and matrix norms, and on $\|\mathbf{B}_Q^{-1}\|_2$. Error tolerances are positive binary rationals of polynomial encoding length. The public bases give the integer-coefficient encoding of ring vectors. These conventions specify the bit-complexity model for our use of Peikert's Gaussian sampler [Pei10].

Theorem 1 (Leaky LWE reduces to Adaptive Leaky LWE)**.** *Fix $\varepsilon_c, \varepsilon_s \in (0, 1/2]$ and $\nu \in (0, 1/2]$ with $\log(1/\varepsilon_s), \log(1/\nu) \in \mathrm{poly}(\lambda)$, and let $\sigma_{f_{\mathsf{LLWE}}} > 0$ be an auxiliary leakage-noise width. Suppose that the following conditions hold:*

(i) Smoothing:
$$\sigma_{f_{\mathsf{LLWE}}} \geq \sqrt{2}\,\eta_{\varepsilon_c}(\Lambda_{n+m}).$$

(ii) Noise correction: *every admissible $\mathbf{L}$ satisfies*
$$\sigma_{f_{\mathsf{LLWE}}}^2 T_{\mathbf{L}} T_{\mathbf{L}}^\top \preceq \frac{1}{2}\Sigma_{\mathbf{f}}.$$

(iii) Gaussian sampling: *for a public parameter r_Q with*
$$r_Q \geq \eta_{\varepsilon_s}(\mathbb{Z}^{\varphi Q}),$$
and a known positive bound $\kappa \geq 2^{-\mathrm{poly}(\lambda)}$,
$$\frac{1}{2}\Sigma_{\mathbf{f}} - r_Q^2 \mathbf{B}_Q \mathbf{B}_Q^\top \succeq \kappa I_{\varphi Q}.$$

Define the Leaky LWE parameter set
$$\mathsf{p}_{\mathrm{id}}^{(t)} := (R, n, m, q, t(n+m), \mathcal{D}_{\mathbf{A}}, \mathcal{D}_{\mathbf{s}}, \mathcal{D}_{\mathbf{e}}, \mathcal{D}_{\mathbf{f}_{\mathsf{LLWE}}}, \{\mathbf{J}_t\}), \tag{2}$$
where its leakage-noise distribution is
$$\mathcal{D}_{\mathbf{f}_{\mathsf{LLWE}}} := D_{R^{t(n+m)}, \sigma_{f_{\mathsf{LLWE}}} I_{\varphi t(n+m)}}, \tag{3}$$
and the sole admissible query is given by
$$\mathbf{J}_t := [\, I_{n+m} \mid \cdots \mid I_{n+m} \,] \in R^{(n+m)\times t(n+m)}.$$

Then, for every admissible PPT adversary $\mathcal{A}$ against $\mathsf{ALLWE}[t]$ with parameters $\mathsf{p} = (R, n, m, q, Q, \mathcal{D}_{\mathbf{A}}, \mathcal{D}_{\mathbf{s}}, \mathcal{D}_{\mathbf{e}}, \mathcal{D}_{\mathbf{f}}, \mathcal{L})$, there is a PPT adversary $\mathcal{R}^{\mathcal{A}}$ against LLWE with parameters $\mathsf{p}_{\mathrm{id}}^{(t)}$ such that
$$\mathrm{Adv}^{\mathsf{LLWE}}_{\mathsf{p}_{\mathrm{id}}^{(t)}, \mathcal{R}^{\mathcal{A}}} \geq \mathrm{Adv}^{\mathsf{ALLWE}[t]}_{\mathsf{p}, \mathcal{A}} - 2t\left(\frac{\varepsilon_c}{1-\varepsilon_c} + 8\varepsilon_s + \nu\right). \tag{4}$$

Remark 1 (A scalar sufficient condition). Condition (ii) holds if
$$\|T_{\mathbf{L}}\|_2^2 \leq \frac{s_{\min}(\Sigma_{\mathbf{f}})}{2\sigma_{f_{\mathsf{LLWE}}}^2} \quad \text{for every admissible } \mathbf{L}.$$
In particular, if $\|T_{\mathbf{L}}\|_2^2 \leq B$ throughout the leakage family, it suffices to choose $s_{\min}(\Sigma_{\mathbf{f}}) \geq 2\sigma_{f_{\mathsf{LLWE}}}^2 B$.

3.2 Proof of the adaptive reduction

For $\mathbf{f}_{\mathsf{LLWE}} \xleftarrow{\$} \mathcal{D}_{\mathbf{f}_{\mathsf{LLWE}}}$, partition $\mathbf{f}_{\mathsf{LLWE}}$ into t consecutive blocks $\mathbf{f}_{\mathsf{LLWE},1}, \ldots, \mathbf{f}_{\mathsf{LLWE},t} \in R^{n+m}$. For each admissible $\mathbf{L}$, define the correction-noise parameter

$$\Sigma_{\mathbf{h}}(\mathbf{L}) := \Sigma_{\mathbf{f}} - \sigma^2_{f_{\mathsf{LLWE}}} T_{\mathbf{L}} T_{\mathbf{L}}^{\top}. \tag{5}$$

The first lemma establishes the distribution of an answer with ideal correction noise; the second gives an efficient sampler for that noise.

Lemma 1 (Gaussian correction)**.** *Assume conditions (i) and (ii). For an admissible* $\mathbf{L}$ *and any* $i \in \{1, \ldots, t\}$*, let*

$$\mathbf{f}_{\mathsf{LLWE},i} \xleftarrow{\$} D_{R^{n+m}, \sigma_{f_{\mathsf{LLWE}}} I_{\varphi(n+m)}}, \quad \mathbf{h} \xleftarrow{\$} D_{R^Q, \sqrt{\Sigma_{\mathbf{h}}(\mathbf{L})}}$$

be independent. Then the distribution of each $\mathbf{L}^{\top}\mathbf{f}_{\mathsf{LLWE},i} + \mathbf{h}$ *is within statistical distance* $\varepsilon_c/(1-\varepsilon_c)$ *of* $\mathcal{D}_{\mathbf{f}} = D_{R^Q, \sqrt{\Sigma_{\mathbf{f}}}}$.

Proof. We apply [GMPW20, Theorem 1]. In real coordinates, the independent vectors $\mathbf{F} := \psi_{n+m}(\mathbf{f}_{\mathsf{LLWE},i})$ and $\mathbf{H} := \psi_Q(\mathbf{h})$ have joint distribution $D_{\Gamma,S}$, where

$$\Gamma := \Lambda_{n+m} \times \Lambda_Q, \quad S := \operatorname{diag}\Big(\sigma_{f_{\mathsf{LLWE}}} I_{\varphi(n+m)}, \sqrt{\Sigma_{\mathbf{h}}(\mathbf{L})}\Big).$$

Let $U := [\, T_{\mathbf{L}} \mid I_{\varphi Q} \,]$ and $J := \binom{I_{\varphi(n+m)}}{-T_{\mathbf{L}}}$. Since $T_{\mathbf{L}}\Lambda_{n+m} \subseteq \Lambda_Q$, we have $U\Gamma = \Lambda_Q$ and $\Gamma \cap \ker U = J\Lambda_{n+m}$, which spans $\ker U$. Moreover, (5) and condition (ii) give $\Sigma_{\mathbf{h}}(\mathbf{L}) \succ 0$ and $\Sigma_{\mathbf{h}}(\mathbf{L}) \succeq \sigma^2_{f_{\mathsf{LLWE}}} T_{\mathbf{L}} T_{\mathbf{L}}^{\top}$, so

$$\|S^{-1}J\|_2^2 = \big\|\sigma^{-2}_{f_{\mathsf{LLWE}}} I_{\varphi(n+m)} + T_{\mathbf{L}}^{\top}\Sigma_{\mathbf{h}}(\mathbf{L})^{-1}T_{\mathbf{L}}\big\|_2 \le 2\sigma^{-2}_{f_{\mathsf{LLWE}}}.$$

Thus the normalized kernel lattice satisfies

$$\eta_{\varepsilon_c}(S^{-1}J\Lambda_{n+m}) \le \|S^{-1}J\|_2\, \eta_{\varepsilon_c}(\Lambda_{n+m}) \le 1$$

by condition (i), with the smoothing parameter computed in the lattice's span. This verifies the theorem's smoothing hypothesis.

The theorem therefore gives a pointwise relative error of at most $2\varepsilon_c/(1-\varepsilon_c)$ between the distribution of $T_{\mathbf{L}}\mathbf{F} + \mathbf{H}$ and $D_{\Lambda_Q, \sqrt{\Sigma_{\mathbf{f}}}}$, since

$$USS^{\top}U^{\top} = \sigma^2_{f_{\mathsf{LLWE}}} T_{\mathbf{L}} T_{\mathbf{L}}^{\top} + \Sigma_{\mathbf{h}}(\mathbf{L}) = \Sigma_{\mathbf{f}}.$$

Summing the absolute probability differences and dividing by two gives statistical distance at most $\varepsilon_c/(1-\varepsilon_c)$. Applying ψ_Q^{-1} proves the claim. □

Lemma 2 (Correction-noise sampler)**.** *For the public parameters and error tolerances of Theorem 1, assume conditions (ii)–(iii) and the numerical representation of Section 3.1. There is a uniform PPT algorithm* Samp *whose output distribution on every admissible input* $\mathbf{L}$ *is within statistical distance* $8\varepsilon_s + \nu$ *of* $D_{R^Q, \sqrt{\Sigma_{\mathbf{h}}(\mathbf{L})}}$.

Proof. Condition (ii) gives $\Sigma_{\mathbf{h}}(\mathbf{L}) \succeq \frac{1}{2}\Sigma_{\mathbf{f}}$. Together with the matrix inequality in condition (iii), this yields

$$\Sigma_{\mathbf{h}}(\mathbf{L}) - r_Q^2 \mathbf{B}_Q \mathbf{B}_Q^{\top} \succeq \kappa I_{\varphi Q}.$$

By condition (iii), $r_Q \ge \eta_{\varepsilon_s}(\mathbb{Z}^{\varphi Q})$, and

$$(r_Q \mathbf{B}_Q)^{-1}\Lambda_Q = r_Q^{-1}\mathbb{Z}^{\varphi Q}.$$

These are the positivity and smoothing conditions for Peikert's generic Gaussian sampler [Pei10, Theorem 2]. With parameter $\Sigma_{\mathbf{h}}(\mathbf{L})$, its ideal output is within $8\varepsilon_s$ of $D_{\Lambda_Q,\sqrt{\Sigma_{\mathbf{h}}(\mathbf{L})}}$.

Under the stated numerical representation and condition (iii), the sampler admits a uniform polynomial-time implementation with total finite-precision and truncation error at most ν. The triangle inequality therefore gives the claimed $8\varepsilon_s + \nu$ bound. □

With these lemmas in place, we are now ready to prove Theorem 1.

Proof of Theorem 1. The adversary $\mathcal{R}^{\mathcal{A}}$ acts as follows.

1. Upon receiving $\mathbf{A}$, it submits $\mathbf{J}_t$ to the Leaky LWE game.
2. Upon receiving $\mathbf{b}$ and the leakage answer, it splits the answer into t blocks $\mathbf{z}_1, \dots, \mathbf{z}_t \in R^{n+m}$. Writing $\mathbf{x} = \binom{\mathbf{s}}{\mathbf{e}}$, these blocks satisfy
$$\mathbf{z}_i = \mathbf{x} + \mathbf{f}_{\mathsf{LLWE},i}, \quad \mathbf{f}_{\mathsf{LLWE},i} \xleftarrow{\$} D_{R^{n+m},\sigma_{f_{\mathsf{LLWE}}} I_{\varphi(n+m)}}, \quad i = 1, \dots, t. \tag{6}$$
The $\mathbf{f}_{\mathsf{LLWE},i}$ are mutually independent and independent of $\mathbf{x}$ and the public data, by the product distribution $\mathcal{D}_{\mathbf{f}_{\mathsf{LLWE}}}$ in (3). The reduction gives $(\mathbf{A}, \mathbf{b})$ to $\mathcal{A}$ and keeps the $\mathbf{z}_i$ private.
3. On the ith valid query $\mathbf{L}_i$, it samples $\widetilde{\mathbf{h}}_i \xleftarrow{\$} \mathsf{Samp}(\mathbf{L}_i)$ using fresh independent randomness and returns
$$\widetilde{\boldsymbol{\ell}}_i := \mathbf{L}_i^\top \mathbf{z}_i + \widetilde{\mathbf{h}}_i = \mathbf{L}_i^\top \mathbf{x} + \mathbf{L}_i^\top \mathbf{f}_{\mathsf{LLWE},i} + \widetilde{\mathbf{h}}_i. \tag{7}$$

The reduction follows the query-validation and output rules of Figure 3 and forwards all oracle queries to $\mathcal{O}_{\mathbf{A}}$ supplied by the Leaky LWE experiment. It runs in polynomial time by Lemma 2 and the polynomial bound on t.

Distribution of each answer. Fix a round i and condition on $\mathbf{x}$, $(\mathbf{A}, \mathbf{b})$, the preceding transcript, the chosen matrix $\mathbf{L}_i$, and the current adversary and oracle states. The unused error $\mathbf{f}_{\mathsf{LLWE},i}$ is independent of these variables and still follows the distribution specified in (6). Lemma 1 bounds the distance between the ideal corrected noise and $\mathcal{D}_{\mathbf{f}}$ by $\varepsilon_c/(1-\varepsilon_c)$. Lemma 2 bounds the additional error from sampling the correction noise by $8\varepsilon_s + \nu$. Hence the conditional distribution of (7) is within

$$\delta := \frac{\varepsilon_c}{1-\varepsilon_c} + 8\varepsilon_s + \nu$$

of a genuine answer $\mathbf{L}_i^\top \mathbf{x} + \mathbf{f}_i$ with fresh $\mathbf{f}_i \xleftarrow{\$} \mathcal{D}_{\mathbf{f}}$.

Advantage loss. A hybrid argument over the t answers gives statistical distance at most $t\delta$ between the simulated and genuine transcripts in each challenge branch. Both experiments use the same rules for invalid queries and termination. Their acceptance probabilities therefore differ by at most $t\delta$ in each branch, yielding the total loss $2t\delta$ in (4). □

3.3 Reduction from LWE

We combine Theorem 1 with [LSW25, Theorem 3] to obtain a reduction from Gaussian-secret LWE to adaptive Leaky LWE.

Corollary 1 (Gaussian-secret LWE reduces to adaptive Leaky LWE)**.** *Let $t = t(\lambda)$ be polynomially bounded, and let $\mathcal{D}_{\mathbf{A}}$ and $\mathcal{L}$ be as in Section 2.2. Denote by the source LWE parameter set*

$$\mathsf{p}^* := (R, n, m, q, \mathcal{D}_{\mathbf{A}}, D_{R^n,\sqrt{\Sigma_{\mathbf{s}}^*}}, D_{R^m,\sqrt{\Sigma_{\mathbf{e}}^*}}),$$

Denote by the target $\mathsf{ALLWE}[t]$ *parameter set*

$$\mathsf{p} := (R, n, m, q, Q, \mathcal{D}_{\mathbf{A}}, D_{R^n,\sqrt{\Sigma_{\mathbf{s}}}}, D_{R^m,\sqrt{\Sigma_{\mathbf{e}}}}, D_{R^Q,\sqrt{\Sigma_{\mathbf{f}}}}, \mathcal{L}).$$

Let ε_0 *be the error tolerance in [LSW25, Theorem 3], and let* $\varepsilon_c, \varepsilon_s, \nu$ *be the tolerances of Theorem 1, all in* $(0, 1/2]$. *Define*

$$\alpha := \left(2\eta_{\varepsilon_0}(\Lambda_{n+m})^2 + \max\{s_{\max}(\Sigma^*_{\mathbf{s}}), s_{\max}(\Sigma^*_{\mathbf{e}})\}\right)^{-1} - \min\{s_{\min}(\Sigma_{\mathbf{s}}), s_{\min}(\Sigma_{\mathbf{e}})\}^{-1}.$$

Suppose that the following conditions hold.

(i) Source and target Gaussian parameters.

$$\min\{s_{\min}(\Sigma^*_{\mathbf{s}}), s_{\min}(\Sigma^*_{\mathbf{e}})\} \geq 2\eta_{\varepsilon_0}(\Lambda_{n+m})^2 \quad \textit{and} \quad \alpha > 0. \tag{8}$$

(ii) Noise correction. *For every admissible* $\mathbf{L}$,

$$\max\left\{\frac{t}{\alpha}, 2\eta_{\varepsilon_c}(\Lambda_{n+m})^2\right\} T_{\mathbf{L}} T_{\mathbf{L}}^\top \preceq \frac{1}{2}\Sigma_{\mathbf{f}}.$$

(iii) Gaussian sampling. *For a known positive bound* $\kappa \geq 2^{-\operatorname{poly}(\lambda)}$,

$$\frac{1}{2}\Sigma_{\mathbf{f}} - \eta_{\varepsilon_s}(\mathbb{Z}^{\varphi Q})^2 \mathbf{B}_Q \mathbf{B}_Q^\top \succeq \kappa I_{\varphi Q},$$

where $\mathbf{B}_Q$ *is the public basis matrix from Section 3.1.*

Then every admissible PPT adversary $\mathcal{A}$ *against* $\mathsf{ALLWE}[t]$ *with parameters* p *yields a PPT LWE distinguisher* $\mathcal{B}$ *with parameters* p^* *such that*

$$\mathrm{Adv}^{\mathsf{LWE}}_{\mathsf{p}^*,\mathcal{B}} \geq \mathrm{Adv}^{\mathsf{ALLWE}[t]}_{\mathsf{p},\mathcal{A}} - \frac{4\varepsilon_0}{1-\varepsilon_0} - 2t\left(\frac{\varepsilon_c}{1-\varepsilon_c} + 8\varepsilon_s + \nu\right). \tag{9}$$

Proof. To apply Theorem 1, take

$$\sigma^2_{f_{\mathsf{LLWE}}} = \max\left\{\frac{t}{\alpha}, 2\eta_{\varepsilon_c}(\Lambda_{n+m})^2\right\}, \quad r_Q = \eta_{\varepsilon_s}(\mathbb{Z}^{\varphi Q}).$$

The choice of $\sigma_{f_{\mathsf{LLWE}}}$ satisfies the smoothing condition, and the matrix inequalities in the corollary give noise correction and Gaussian sampling.

Theorem 1 gives an adversary $\mathcal{R}^{\mathcal{A}}$ for $\mathsf{p}^{(t)}_{\mathrm{id}}$ with the loss in (4). Equation (8) gives the source-width requirement of [LSW25, Condition 1]. The map $T_{\mathbf{J}_t}$ stacks t copies of the identity, so

$$T_{\mathbf{J}_t}^\top T_{\mathbf{J}_t} = tI_{\varphi(n+m)} \quad \text{and} \quad \|T_{\mathbf{J}_t}\|_2^2 = t \leq \sigma^2_{f_{\mathsf{LLWE}}}\alpha.$$

Since the leakage-noise parameter matrix in (3) has smallest eigenvalue $\sigma^2_{f_{\mathsf{LLWE}}}$, $\mathbf{J}_t$ is admissible under [LSW25, Condition 1]. Applying [LSW25, Theorem 3] costs $4\varepsilon_0/(1-\varepsilon_0)$ in advantage. Adding the loss in (4) gives (9). □

For polynomially bounded t, the total loss is negligible when $\varepsilon_0, \varepsilon_c, \varepsilon_s$ and ν are negligible.

4 Explicit Parameters and Special Cases

We first give explicit Gaussian-width and leakage-norm bounds for power-of-two cyclotomic rings. We then describe how to choose the leakage matrix to obtain several common hint forms.

4.1 Explicit Parameter Conditions

We specialize Corollary 1 to Gaussian-secret Module-LWE over power-of-two cyclotomic rings. The following conditions give explicit bounds on the source, target, and leakage widths for polynomially many adaptive leakage queries.

Corollary 2 (Gaussian-secret Module-LWE reduces to adaptive Leaky LWE)**.** *Let $R = \mathbb{Z}[X]/(X^N+1)$, where N is a power of two, and use the coefficient embedding as in Section 2.1, so that $\varphi = N$ and $\Lambda_k = \mathbb{Z}^{Nk}$. Let $\mathcal{D}_{\mathbf{A}}$ be the distribution that samples $\mathbf{A} \xleftarrow{\$} U(R_q^{n\times m})$ with an empty oracle $\mathcal{O}_{\mathbf{A}}$. Denote by the source Module-LWE parameter set*

$$\mathsf{p}^* := (R, n, m, q, \mathcal{D}_{\mathbf{A}}, D_{R^n,\sigma_s^* I_{Nn}}, D_{R^m,\sigma_e^* I_{Nm}}),$$

Let $t \geq 1$ be polynomially bounded, and let $\mathcal{L}$ be a polynomial-time decidable leakage family. Let $B \geq 0$ be a public bound such that $\|T_{\mathbf{L}}\|_2^2 \leq B$ for every admissible $\mathbf{L}$. Denote by the target $\mathsf{ALLWE}[t]$ parameter set

$$\mathsf{p} := (R, n, m, q, Q, \mathcal{D}_{\mathbf{A}}, D_{R^n,\sigma_s I_{Nn}}, D_{R^m,\sigma_e I_{Nm}}, D_{R^Q,\sigma_f I_{NQ}}, \mathcal{L}).$$

Fix $\varepsilon \in (0, 1/2]$ with $\log(1/\varepsilon) \in \mathrm{poly}(\lambda)$ and a known positive parameter $\kappa \geq 2^{-\mathrm{poly}(\lambda)}$. Define

$$\underline{\alpha} := \big(2\eta_\varepsilon(\mathbb{Z}^N)^2 + \max\{(\sigma_s^*)^2, (\sigma_e^*)^2\}\big)^{-1} - \min\{\sigma_s^2, \sigma_e^2\}^{-1}.$$

Suppose that

$$\begin{aligned} \min\{\sigma_s^*, \sigma_e^*\} &\geq \sqrt{2}\,\eta_\varepsilon(\mathbb{Z}^N), \quad \underline{\alpha} > 0, \\ \sigma_f &\geq \sqrt{2}\max\left\{\sqrt{\frac{tB}{\underline{\alpha}}},\ \eta_\varepsilon(\mathbb{Z}^{NQ}) + \kappa\right\}. \end{aligned} \tag{10}$$

Then every admissible PPT adversary $\mathcal{A}$ against $\mathsf{ALLWE}[t]$ with target parameters p yields a PPT distinguisher $\mathcal{B}$ for Gaussian-secret Module-LWE with source parameters p^* such that

$$\mathrm{Adv}^{\mathsf{MLWE}}_{\mathsf{p}^*,\mathcal{B}} \geq \mathrm{Adv}^{\mathsf{ALLWE}[t]}_{\mathsf{p},\mathcal{A}} - \big(8(n+m) + 22t\big)\varepsilon.$$

Proof. For $B < 1/2$, every admissible query is zero because $T_{\mathbf{L}}$ has integer entries, so $\mathsf{ALLWE}[t]$ is equivalent to MLWE at the target widths. By (10), the source widths and the added Gaussian widths $\sqrt{\sigma_u^2 - (\sigma_u^*)^2}$ for $u \in \{s, e\}$ are at least $\sqrt{2}\,\eta_\varepsilon(\mathbb{Z}^N)$, so [LSW25, Lemma 7] converts the widths ring coordinate by ring coordinate with loss at most $8(n+m)\varepsilon$.

For $B \geq 1/2$, assume $\big(8(n+m) + 22t\big)\varepsilon < 1$, since otherwise the claimed advantage lower bound is at most zero. Apply Corollary 1 with

$$\varepsilon_0 = (1+\varepsilon)^{n+m} - 1, \quad \varepsilon_c = \varepsilon_s = \nu = \varepsilon.$$

Since $(n+m)\varepsilon < 1/8$, we have $\varepsilon_0 < 1/2$ and $\varepsilon_0/(1-\varepsilon_0) \leq 2(n+m)\varepsilon$. Multiplying the $n+m$ Gaussian sums on $\mathbb{Z}^N$ gives $\eta_{\varepsilon_0}(\Lambda_{n+m}) = \eta_\varepsilon(\mathbb{Z}^N)$, hence $\alpha = \underline{\alpha}$ and (8) follows from (10).

Since $(1+\varepsilon^2)^{n+m} - 1 \leq \varepsilon$, the definition of $\underline{\alpha}$ and the source-width bound give

$$2\eta_\varepsilon(\Lambda_{n+m})^2 \leq 4\eta_\varepsilon(\mathbb{Z}^N)^2 < \frac{1}{\underline{\alpha}} \leq \frac{t}{\underline{\alpha}}.$$

The maximum in condition (ii) is therefore $t/\underline{\alpha}$. With $\mathbf{B}_Q = I_{NQ}$ and $\eta_\varepsilon(\mathbb{Z}^{NQ}) > 1/2$, (10) gives noise correction and Gaussian sampling

$$\frac{t}{\underline{\alpha}} T_{\mathbf{L}} T_{\mathbf{L}}^\top \preceq \frac{tB}{\underline{\alpha}} I_{NQ} \preceq \frac{1}{2}\Sigma_{\mathbf{f}}, \quad \left(\frac{\sigma_f^2}{2} - \eta_\varepsilon(\mathbb{Z}^{NQ})^2\right) I_{NQ} \succeq \kappa I_{NQ}.$$

Since $\varepsilon/(1-\varepsilon) \leq 2\varepsilon$, (9) gives the claimed loss. □

4.2 Specializations of the Leakage Matrix

After fixing the ring, Gaussian distributions, and hint dimension, one may restrict the admissible leakage family without changing the reduction or its loss. Indeed, every polynomial-time decidable subfamily $\mathcal{L}' \subseteq \mathcal{L}$ inherits the operator bound and uniform sampling guarantee for $\mathcal{L}$. Corollary 1 therefore applies with $\mathcal{L}'$ in place of $\mathcal{L}$. The explicit parameter conditions above provide sufficient choices when the leakage noise is spherical.

For error-only leakage, take a polynomial-time decidable family $\mathcal{L}_e \subseteq R^{m\times Q}$ and set $\mathcal{L} = \{(\mathbf{0}, \mathbf{L}_e) : \mathbf{L}_e \in \mathcal{L}_e\}$. For secret-only leakage, take $\mathcal{L}_s \subseteq R^{n\times Q}$ and set $\mathcal{L} = \{(\mathbf{L}_s, \mathbf{0}) : \mathbf{L}_s \in \mathcal{L}_s\}$. The respective answers are

$$\boldsymbol{\ell}_i = \mathbf{L}_{e,i}^\top \mathbf{e} + \mathbf{f}_i, \quad \boldsymbol{\ell}_i = \mathbf{L}_{s,i}^\top \mathbf{s} + \mathbf{f}_i.$$

In either case, the nonzero matrix may depend on $(\mathbf{A}, \mathbf{b})$ and all preceding answers. These are adaptive versions of the corresponding leakage forms, under the Gaussian and sampling conditions of our reduction. The error-only case has the leakage form of Error-Leakage LWE [DKL+23] and still requires a Gaussian secret, as specified in Corollary 1.

For the structured hints of Hint-MLWE [KLSS23], let r be polynomially bounded, let $\Gamma \subseteq R^r$ be polynomial-time decidable, and define

$$\mathbf{L}_{\boldsymbol{\gamma}} := [\,\gamma_1 I_{n+m} \mid \cdots \mid \gamma_r I_{n+m}\,] \in R^{(n+m)\times r(n+m)}, \quad \boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_r) \in \Gamma.$$

Set $Q = r(n+m)$ and let the leakage family consist of the first n and last m rows of these matrices. An answer has the form

$$\boldsymbol{\ell} = \mathbf{L}_{\boldsymbol{\gamma}}^\top \mathbf{x} + \mathbf{f}.$$

Our adaptive experiment allows a new tuple $\boldsymbol{\gamma}$ to be selected after the challenge and preceding answers. Fixing the tuple, or sampling it independently from a prescribed public joint distribution supported on Γ, recovers the corresponding restricted choices. The reduction applies when the resulting matrices meet the operator and sampling conditions.

The rectangular hints of [WSS+26] can be written in the same notation as

$$\boldsymbol{\ell}_i = \mathbf{L}_i^\top \mathbf{x} + \mathbf{f}_i, \quad \mathbf{L}_i \in R^{(n+m)\times Q}.$$

Their notation $\boldsymbol{\ell}_i = \mathbf{R}_i \mathbf{x} + \mathbf{f}_i$ corresponds to $\mathbf{L}_i = \mathbf{R}_i^\top$. If the public matrix is sampled from a prescribed distribution, an adversary in our experiment can make the same query by sampling and submitting that matrix, provided its support lies in a polynomial-time decidable admissible family. Allowing the matrix instead to depend on the challenge and earlier answers gives the adaptive extension under our parameter conditions.

Finally, setting $Q = 1$ makes each query a vector $\mathbf{L}_i \in R^{(n+m)\times 1}$ and each answer a single ring element

$$\ell_i = \mathbf{L}_i^\top \mathbf{x} + f_i \in R.$$

In the uniform-public-matrix cyclotomic setting with spherical Gaussians, this is the scalar-hint setting compared with the concurrent Adaptive Hint-MLWE work in Section 5.

5 Comparison with Concurrent Work

We compare the leakage-noise width required by Corollary 2 and by [HKM+26, Theorem 1] with the same query count, query-norm bound, and advantage-loss budget. We use the coefficient embedding of $R = \mathbb{Z}[X]/(X^N + 1)$, where N is a power of two, and take a uniform public matrix $\mathbf{A}$ with an empty oracle $\mathcal{O}_{\mathbf{A}}$. The adversary makes $t \geq 1$ adaptive queries, with t polynomially bounded. Each answer is one element of R in [HKM+26, Theorem 1], so $Q = 1$ in Corollary 2. The source secret and error have width σ^*, the

target secret and error have width σ, and the leakage noise has width σ_f. We compare σ_f/σ^* using the Gaussian convention as stated in Table 1. Throughout this section, "their" refers to Hwang et al. [HKM+26].

For each query $\mathbf{L}_i = (L_{i,j})_{j=1}^{n+m} \in R^{(n+m)\times 1}$, require

$$\sum_{j=1}^{n+m} \|L_{i,j}\|_1^2 \le B, \quad B \ge 1,$$

where $\|\cdot\|_1$ is the sum of the absolute values of the coefficients of a ring element. We use this B in Corollary 2 because $\|T_{\mathbf{L}_i}\|_2^2 \le \sum_{j=1}^{n+m} \|L_{i,j}\|_1^2$. In [HKM+26, Theorem 1], we take $B_{\text{hint}} = B$.

Fix an advantage-loss budget $\delta \in (0,1)$ with $\log(1/\delta) \in \text{poly}(\lambda)$ and set

$$\varepsilon = \frac{\delta}{8(n+m)+22t}, \quad \varepsilon' = \frac{\delta}{4t+2(n+m)}.$$

Both advantage-loss bounds equal δ. Since $\varepsilon \le \varepsilon'$, we have $\eta_{\varepsilon'}(\mathbb{Z}^N) \le \eta_\varepsilon(\mathbb{Z}^N)$. The resulting conditions on σ^* and σ are as follows.

	Our reduction	[HKM+26]
Source width	$\sigma^* \ge \sqrt{2}\,\eta_\varepsilon(\mathbb{Z}^N)$	$\sigma^* \ge \sqrt{2}\,\eta_{\varepsilon'}(\mathbb{Z}^N)$
Target-to-source width ratio	$\dfrac{\sigma}{\sigma^*} > \sqrt{1+\dfrac{2\eta_\varepsilon(\mathbb{Z}^N)^2}{(\sigma^*)^2}}$	$\dfrac{\sigma}{\sigma^*} > \sqrt{2}$

In (10), take $Q = 1$ and choose $\kappa = \eta_\varepsilon(\mathbb{Z}^N)$. The two terms in the maximum are then $\sqrt{tB/\underline{\alpha}}$ and $2\eta_\varepsilon(\mathbb{Z}^N)$. By the definition of $\underline{\alpha}$, the bound on σ^*, and $B, t \ge 1$,

$$\frac{tB}{\underline{\alpha}} \ge \frac{1}{\underline{\alpha}} > (\sigma^*)^2 + 2\eta_\varepsilon(\mathbb{Z}^N)^2 \ge 4\eta_\varepsilon(\mathbb{Z}^N)^2.$$

Thus the first term is larger, and (10) requires $\sigma_f \ge \sqrt{2tB/\underline{\alpha}}$. Substituting the definition of $\underline{\alpha}$ gives

$$\frac{\sigma_f}{\sigma^*} \ge \sqrt{\frac{2Bt}{\left(1+2\eta_\varepsilon(\mathbb{Z}^N)^2/(\sigma^*)^2\right)^{-1} - (\sigma^*/\sigma)^2}}. \tag{11}$$

Theorem 1 of [HKM+26] uses Q for the number of queries. We set their Q to our t, substitute their formula for κ, and multiply their Gaussian widths by $\sqrt{2\pi}$ to use our convention. In this way, their equality becomes

$$\frac{1}{(\sigma^*)^2} - \frac{2}{\sigma^2} = \frac{t}{\eta_{\varepsilon'}(\mathbb{Z}^N)^2}\left(1 - \sqrt{1 - \frac{4B\eta_{\varepsilon'}(\mathbb{Z}^N)^2}{\sigma_f^2}}\right).$$

Solving for σ_f/σ^* gives

$$\frac{\sigma_f}{\sigma^*} = \sqrt{\frac{2Bt}{1 - 2(\sigma^*/\sigma)^2 - \dfrac{\eta_{\varepsilon'}(\mathbb{Z}^N)^2}{2t(\sigma^*)^2}\left(1-2(\sigma^*/\sigma)^2\right)^2}}. \tag{12}$$

For a fixed target-to-source width ratio $\sigma/\sigma^* > \sqrt{2}$, both bounds scale as $\sqrt{Bt}$ when $\eta_\varepsilon(\mathbb{Z}^N)/\sigma^*$ tends to zero. For example, at $\sigma = 2\sigma^*$, the leading constants are $\sqrt{8/3}$ for our bound (11) and 2 for their bound (12). Our sufficient leakage width is smaller in this regime. At our minimum source width $\sigma^* = \sqrt{2}\,\eta_\varepsilon(\mathbb{Z}^N)$, their bound is smaller for the same target-to-source width ratio.

Acknowledgement

This work is partly supported by Australian Research Council Discovery Grants DP250100229 and DP260100245.

References

[ALS16] Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In *Annual International Cryptology Conference*, pages 333–362. Springer, 2016. doi:10.1007/978-3-662-53015-3_12.

[ASP12] Jacob Alperin-Sheriff and Chris Peikert. Circular and KDM security for identity-based encryption. In *International Workshop on Public Key Cryptography*, pages 334–352. Springer, 2012. doi:10.1007/978-3-642-30057-8_20.

[BGV14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)*, 6(3):1–36, 2014. doi:10.1145/2633600.

[BJTW26] Katharina Boudgoust, Corentin Jeudy, Erkan Tairi, and Weiqiang Wen. Hardness of M-LWE with general distributions and applications to leaky variants. In *IACR International Conference on Public-Key Cryptography*, pages 3–37. Springer, 2026. doi:10.1007/978-3-032-26731-3_1.

[BKL+25] Cecilia Boschini, Darya Kaviani, Russell W. F. Lai, Giulio Malavolta, Akira Takahashi, and Mehdi Tibouchi. Ringtail: Practical two-round threshold signatures from learning with errors. In *2025 IEEE Symposium on Security and Privacy*, pages 149–164. IEEE, 2025. doi:10.1109/SP61157.2025.00070.

[DKL+23] Nico Döttling, Dimitris Kolonelos, Russell W. F. Lai, Chuanwei Lin, Giulio Malavolta, and Ahmadreza Rahimi. Efficient laconic cryptography from learning with errors. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 417–446. Springer, 2023. doi:10.1007/978-3-031-30620-4_14.

[DPKM+24] Rafaël Del Pino, Shuichi Katsumata, Mary Maller, Fabrice Mouhartem, Thomas Prest, and Markku-Juhani Saarinen. Threshold Raccoon: practical threshold signatures from standard lattice assumptions. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 219–248. Springer, 2024. doi:10.1007/978-3-031-58723-8_8.

[ENP24] Thomas Espitau, Guilhem Niot, and Thomas Prest. Flood and submerse: Distributed key generation and robust threshold signature from lattices. In *Annual International Cryptology Conference*, pages 425–458. Springer, 2024. doi:10.1007/978-3-031-68394-7_14.

[ESTX26] Muhammed F Esgin, Ron Steinfeld, Erkan Tairi, and Jie Xu. LeOPaRd: Towards practical post-quantum oblivious PRFs via 2HashDH paradigm. In *Annual International Cryptology Conference*, pages 35–67. Springer, 2026. doi:10.1007/978-3-032-35377-1_2.

[GMPW20] Nicholas Genise, Daniele Micciancio, Chris Peikert, and Michael Walter. Improved discrete Gaussian and Subgaussian analysis for lattice cryptography. In *IACR International Conference on Public-Key Cryptography*, pages 623–651. Springer, 2020. doi:10.1007/978-3-030-45374-9_21.

[HKM+26] Yisol Hwang, Shuichi Katsumata, Seonhong Min, Guilhem Niot, and Yongsoo Song. Proving threshold Regev PKE from adaptive Hint-MLWE: Efficient, non-interactive, and CCA secure. *Cryptology ePrint Archive*, 2026. URL: https://eprint.iacr.org/2026/1585.

[KLSS23] Duhyeong Kim, Dongwon Lee, Jinyeong Seo, and Yongsoo Song. Toward practical lattice-based proof of knowledge from Hint-MLWE. In *Annual International Cryptology Conference*, pages 549–580. Springer, 2023. doi:10.1007/978-3-031-38554-4_18.

[KSSY26] Yuki Kume, Ron Steinfeld, Amin Sakzad, and Mert Yassi. LUNA+: More succinct post-quantum ZK-SNARKs from computational privacy. *Cryptology ePrint Archive*, 2026. To appear in CCS 2026. URL: https://eprint.iacr.org/2026/1639.

[LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *Journal of the ACM (JACM)*, 60(6):1–35, 2013. doi:10.1145/2535925.

[LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, 2015. doi:10.1007/s10623-014-9938-4.

[LSW25] Russell W. F. Lai, Monisha Swarnakar, and Ivy K. Y. Woo. Leaky LWE: Learning with errors with semi-adaptive secret- and error-leakage. *IACR Communications in Cryptology*, 2(3):1–23, 2025. doi:10.62056/ah89ksuc2.

[Pei10] Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In *Annual Cryptology Conference*, pages 80–97. Springer, 2010. doi:10.1007/978-3-642-14623-7_5.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, 2009. doi:10.1145/1568318.1568324.

[SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 617–635. Springer, 2009. doi:10.1007/978-3-642-10366-7_36.

[WLL24] Zhedong Wang, Qiqi Lai, and Feng-Hao Liu. Ring/module learning with errors under linear leakage–hardness and applications. In *IACR International Conference on Public-Key Cryptography*, pages 275–304. Springer, 2024. doi:10.1007/978-3-031-57722-2_9.

[WSS+26] Hongxiao Wang, Ron Steinfeld, Markku-Juhani O Saarinen, Muhammed F Esgin, and Siu-Ming Yiu. mmCipher: Batching post-quantum public key encryption made bandwidth-optimal. In *35th USENIX Security Symposium, USENIX Security*, pages 4601–4620, 2026. doi:10.5555/3857392.3857625.