LaMS: A p-adic Layered Modulus Switching for Provable Dual Attacks on LWE

Rui-Jie Wang, Zhong-Xiao Wang(✉), Qun-Xiong Zheng(✉), and Xuan Zhao

Information Engineering University, Zhengzhou 450001, China
ruijie_wang1@163.com, zhongxiao_wang@126.com
qunxiong_zheng@163.com, xuanzhao5280@outlook.com

Abstract. The Learning with Errors (LWE) problem is a central foundation for post-quantum schemes such as Kyber and Dilithium. Dual attacks are among the main tools for assessing the concrete hardness of LWE instances. At EUROCRYPT 2024, Pouly and Shen introduced the first provable dual attack against LWE. Subsequently, at ASIACRYPT 2025, Qu and Xu incorporated modulus switching into this framework by recovering the guessed secret mod several small primes and recombining the resulting residues via the Chinese Remainder Theorem (CRT). Although this CRT-based strategy substantially reduces the search space, it reconstructs the full guessed secret through several distinct primes whose product must exceed q. Consequently, the total guessing cost is dominated by the largest CRT prime p_k. This raises a natural question: can the same recovery effect be achieved by repeatedly applying the subroutine with a fixed small prime, while further reducing the overall complexity?

We answer this question affirmatively by proposing layered modulus switching (LaMS), a provable modulus switching dual attack based on a p-adic view of the guessed secret. Instead of recovering residues mod several distinct primes, LaMS fixes a single small prime p and recovers the guessed secret digit by digit in its p-adic expansion. After each digit is recovered, its contribution is subtracted from the LWE samples, producing a new target LWE instance in which the next digit becomes the new secret mod p. As a result, the guessing cost is reduced from $O(\mathrm{poly}(m,n)(N+\sum_{j=1}^{k} p_j^{n_{\mathrm{guess}}}))$ in the CRT-based attack to $O(\mathrm{poly}(m,n)\lceil \log_p q \rceil (N+p^{n_{\mathrm{guess}}}))$, where $p \ll p_k$.

We also correct a parameter issue in two previous works (EUROCRYPT 2024 and ASIACRYPT 2025) on Kyber estimates. After this correction, LaMS lowers the estimated attack cost by up to 1 bit for Kyber compared with the corrected CRT-based attack of Qu and Xu.

Keywords: LWE · Provable dual attack · Modulus switching · CRT · LaMS

1 Introduction

The Learning with Errors (LWE) problem [17], introduced by Regev in 2005, has become a cornerstone of post-quantum cryptography owing to its quantum-

resistant hardness properties. It underpins the security of widely deployed post-quantum schemes such as Kyber [18] and Dilithium [10], both of which have been standardized by NIST. The concrete security of these schemes depends critically on accurate hardness estimates for LWE, making the analysis of LWE attacks a central research topic.

Two classical families of attacks against LWE are primal attacks and dual attacks. This paper focuses primarily on the latter. Dual attacks aim to find many short vectors in a suitable dual lattice and use them to distinguish LWE samples from uniformly random samples. The distinguishing signal arises from the fact that inner products with short dual vectors preserve a detectable bias when evaluated on LWE samples, whereas the corresponding values are uniformly distributed for truly random samples.

A line of work has refined this idea into increasingly efficient concrete attacks. Albrecht [2] split the secret into a guessed part and a lattice part, using lattice reduction to construct short dual vectors for the distinguishing step. Espitau et al. [8] accelerated the guessing phase by exploiting the recursive structure of the search space, and Guo and Johansson [9] further improved the attack by applying the Fast Fourier Transform (FFT). The MATZOV report [11] introduced modulus switching into this framework, reducing the FFT domain from $\mathbb{Z}_q^{n_{\text{fft}}}$ to $\mathbb{Z}_p^{n_{\text{fft}}}$ for a much smaller prime $p < q$, thereby yielding substantial practical speedups. At CRYPTO 2025, Carrier et al. [6] replaced modulus switching with an efficient polar-code decoding algorithm and provided complexity estimates for Kyber that fall below the NIST security requirements.

Despite these advances, Ducas and Pulles [7] showed that the independent score heuristic underlying the dual attacks of [9,11] can fail dramatically. In particular, they identified a contradictory regime in which this heuristic no longer provides a reliable prediction. This motivated the search for dual attacks whose success can be justified without relying on the problematic independent score heuristic.

Pouly and Shen [14] addressed this challenge by proposing the first provable dual attack on LWE at EUROCRYPT 2024. Their work replaces the heuristic score analysis with rigorous estimates from the geometry of lattices. They also define an effective regime in which the attack is guaranteed to succeed, and prove that this regime is disjoint from the contradictory regime of [7]. Nevertheless, their guessing phase still enumerates the full space $\mathbb{Z}_q^{n_{\text{guess}}}$, leaving a major bottleneck in its concrete complexity.

Qu and Xu [16] subsequently incorporated modulus switching into the provable dual framework. Their main idea is to replace the guessing matrix by a rounded matrix mod a smaller prime p, reducing the enumeration space from $\mathbb{Z}_q^{n_{\text{guess}}}$ to $\mathbb{Z}_p^{n_{\text{guess}}}$. The rounding residual is embedded into a rational q-ary lattice, so the procedure does not introduce an additional error term. A single run of this subroutine recovers only $\mathbf{s}_{\text{guess}} \bmod p$, rather than the full guessed secret. Therefore, to reconstruct $\mathbf{s}_{\text{guess}}$, the attack is repeated for several distinct primes $p_1 < p_2 < \cdots < p_k$ satisfying $\prod_j p_j > q$, and the resulting residues are combined via the Chinese Remainder Theorem (CRT). This CRT-based strategy

reduces the search domain dramatically, yielding a 15–29 bit improvement over the provable dual attack of [14] [1].

In this paper, we revisit the CRT-based modulus switching strategy of [16]. Although this strategy substantially reduces the enumeration space by working mod small primes, it reconstructs the guessed secret through several independent runs mod distinct primes, and the total guessing cost is ultimately dominated by the largest CRT prime. This motivates the central question of our work: can one recover the full guessed secret by repeatedly applying the subroutine with a fixed small prime, thereby avoiding larger CRT primes and further reducing the overall complexity? To answer this question, we fix a single small prime p, viewing the guessed secret through its p-adic expansion, and recovering its digits layer by layer.

1.1 Our Contribution

The contributions of this work are twofold.

1. **A p-adic Layered Modulus Switching Dual Attack.** The key idea of LaMS is to fix a single small prime $p < q$, represent the guessed secret in its p-adic expansion, and recover its digits successively, while all digit-recovery steps are performed mod the same prime p. Consequently, all layers work over the same search space $\mathbb{Z}_p^{n_{\text{guess}}}$. In terms of complexity, LaMS replaces the guessing cost $O(\text{poly}(m,n)(N + \sum_{j=1}^{k} p_j^{n_{\text{guess}}}))$ in the CRT-based attack of [16] with $O(\text{poly}(m,n)\lceil \log_p q \rceil (N + p^{n_{\text{guess}}}))$. Therefore, whenever the fixed prime p is substantially smaller than the largest CRT prime p_k, the layered strategy yields a concrete complexity improvement.
2. **Corrected Complexity Estimates and Concrete Improvements.** During our re-evaluation of the concrete complexities, we found that the estimates of [14,16] used a slightly inconsistent error-parameter substitution. We apply this correction throughout our complexity analysis and re-estimate the concrete bit-security of Kyber for both the CRT-based modulus switching dual attack of [16] and our LaMS. The resulting comparison shows that LaMS improves over the CRT-based attack by up to 1 bit for Kyber.

1.2 Organization

Section 2 introduces the notation and preliminary background. Section 3 reviews the provable dual attack of [14] and the CRT-based modulus switching of [16]. Section 4 presents the LaMS attack and gives the concrete complexity estimates. Appendix B compares the effective regime of LaMS with the contradictory regime of [7]. Section 5 concludes the paper.

[1] This figure is based on the original complexity estimates reported in [16], which have since been updated in the latest ePrint version [15]. Beyond this revision, we identified an additional error and provide the latest corrected estimates in Table 1 below.

2 Preliminaries

This section collects the notation and related concepts used throughout the paper.

2.1 Notation and Parameter Conventions

Let $\mathbb{N}$, $\mathbb{R}$, and $\mathbb{Z}$ denote the sets of natural numbers, real numbers, and integers, respectively. For an integer $q \geq 2$, the set $\mathbb{Z}_q = \{0, 1, \ldots, q-1\}$ is used to denote the residue classes mod q, represented by their canonical representatives. Bold lowercase letters denote vectors and bold uppercase letters denote matrices. Unless stated otherwise, all vectors are column vectors. The symbols log and ln denote the base-2 logarithm and the natural logarithm, respectively.

For $x \in \mathbb{R}$, let $\lfloor x \rfloor$ be the floor function and let $\{x\} = x - \lfloor x \rfloor$. These operations are applied entry-wise to vectors and matrices. For vectors $\mathbf{x}$ and $\mathbf{y}$, let $\langle \mathbf{x}, \mathbf{y} \rangle$ be the standard inner product and $\|\mathbf{x}\|$ the Euclidean norm. For a finite set X, $U(X)$ denotes the uniform distribution on X, and $x \leftarrow \chi$ means that x is sampled from the distribution χ. The total variation distance between two distributions P and Q on countable X is denoted as $d_X(P, Q) = \frac{1}{2} \sum_{x \in X} |P(x) - Q(x)|$.

To avoid confusion between the index i and the imaginary unit, we denote the latter by i.

2.2 Learning With Errors

We first recall some basic concepts related to LWE that will be used throughout this paper.

Definition 1 ([17]). *Let q be a prime and n, m be positive integers, let χ be a probability distribution on $\mathbb{Z}$ and let $\mathbf{s} \in \mathbb{Z}_q^n$ be a secret vector. A pair*

$$(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \mod q) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$$

is called an LWE instance, where $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is chosen uniformly at random and $\mathbf{e} \leftarrow \chi^m$. Here, χ^m means that each component of $\mathbf{e}$ is distributed according to χ.

The LWE problem has two versions:

- **Decision-LWE**: Given a pair $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$, which is drawn either from the LWE distribution $(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \mod q)$ or from the uniform distribution on $\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$, determine which distribution it comes from.
- **Search-LWE**: Given an LWE instance $(\mathbf{A}, \mathbf{b} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \mod q) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$, recover the secret vector $\mathbf{s}$.

This paper studies the search-LWE problem. We do not make any assumption on the distribution of the secret. The error distribution χ is typically a discrete Gaussian or a centered binomial distribution in lattice-based cryptographic schemes.

2.3 The q-ary and Rational q-ary Lattices

We first recall some basic concepts from lattice.

Definition 2. *A lattice is a discrete additive subgroup of $\mathbb{R}^m$. Specifically, given n linearly independent vectors $\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n \in \mathbb{R}^m$, the set*

$$\mathcal{L} = \{\mathbf{v} : \mathbf{v} = c_1\mathbf{b}_1 + c_2\mathbf{b}_2 + \cdots + c_n\mathbf{b}_n, c_1, c_2, \ldots, c_n \in \mathbb{Z}\}$$

is called an m-dimensional lattice of rank n. The n linearly independent vectors $\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n$ are called a basis of the lattice $\mathcal{L}$. A lattice $\mathcal{L}$ is called a full-rank lattice if $n = m$.

Given $\mathbf{x} \in \mathbb{R}^m$, its distance to $\mathcal{L}$ is $\mathrm{dist}(\mathbf{x}, \mathcal{L}) = \min_{\mathbf{v}\in\mathcal{L}} \|\mathbf{x} - \mathbf{v}\|$. Let $\mathbf{B}$ be the $m \times n$ matrix consisting of column vectors $\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n$. The determinant of the lattice $\mathcal{L}$ is defined by $\det(\mathcal{L}) = \sqrt{\det(\mathbf{B}^\top\mathbf{B})}$, where $\mathbf{B}^\top$ is the transpose of $\mathbf{B}$.

Let $\mathcal{L}$ be a full-rank lattice in $\mathbb{R}^m$, and let $S \subseteq \mathbb{R}^m$ be a measurable set with volume $\mathrm{vol}(S)$. The Gaussian heuristic [12] predicts that $\mathcal{L} \cap S$ contains roughly $\mathrm{vol}(S)/\det(\mathcal{L})$ lattice points. In particular, the Gaussian heuristic estimate $\mathrm{GH}(\mathcal{L}) = \sqrt{\frac{m}{2\pi e}}\ \det(\mathcal{L})^{1/m}$ is commonly used as an estimate for $\lambda_1(\mathcal{L})$, the length of a shortest nonzero vector in $\mathcal{L}$.

The dual lattice of $\mathcal{L}$ is

$$\mathcal{L}^\vee \ = \ \{\mathbf{u} \in \mathrm{span}(\mathcal{L}) : \langle \mathbf{u}, \mathbf{v}\rangle \in \mathbb{Z} \text{ for all } \mathbf{v} \in \mathcal{L}\},$$

where $\mathrm{span}(\mathcal{L})$ denotes the linear subspace over $\mathbb{R}$ spanned by the vectors in $\mathcal{L}$.

Given a matrix $\mathbf{C} \in \mathbb{Z}_q^{m\times k}$, the ordinary q-ary lattices used in the dual attack are

$$\mathcal{L}_q(\mathbf{C}) = \left\{\mathbf{x} \in \mathbb{Z}^m : \exists\, \mathbf{s} \in \mathbb{Z}^k \text{ such that } \mathbf{C}\mathbf{s} \equiv \mathbf{x} \bmod q\right\},$$
$$\mathcal{L}_q^\perp(\mathbf{C}) = \left\{\mathbf{x} \in \mathbb{Z}^m : \mathbf{C}^\top\mathbf{x} \equiv 0 \bmod q\right\}.$$

It is well-known that $\det(\mathcal{L}_q(\mathbf{C})) = q^{m-r}$ and $\det(\mathcal{L}_q^\perp(\mathbf{C})) = q^r$ if $\mathbf{C}$ has rank r over $\mathbb{Z}_q$. Moreover, $\left(\mathcal{L}_q(\mathbf{C})\right)^\vee = \frac{1}{q}\mathcal{L}_q^\perp(\mathbf{C})$, which can be verified directly.

Modulus switching is a useful tool in dual attacks, as it can reduce the attack complexity. However, it naturally introduces an additional rounding error. In [16], this error is not treated as an additional noise term. Instead, the corresponding additional error is embedded into a rational lattice, leading to the notion of rational q-ary lattices. We recall its formal definition below.

Definition 3 ([16, Definition 1]). *Let $p < q$ be distinct primes, let $m, n \in \mathbb{N}$ with $m \geq n$, and let $0 \leq k \leq n$. Let $\mathbf{A}_1 \in \left(\frac{1}{p}\mathbb{Z}_q\right)^{m\times k}$, $\mathbf{A}_2 \in \mathbb{Z}_q^{m\times(n-k)}$, and set $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$, where $\frac{1}{p}\mathbb{Z}_q = \{0, \frac{1}{p}, \frac{2}{p}, ..., \frac{q-1}{p}\}$. The lattice $\mathcal{L}_q(\mathbf{A})$ is called the rational q-ary lattice associated with $\mathbf{A}$.*

It follows directly from the definition that $q\mathbb{Z}^m \subseteq \mathcal{L}_q(\mathbf{A}) \subseteq \frac{1}{p}\mathbb{Z}^m$ and $pq\mathbb{Z}^m \subseteq \mathcal{L}_q^{\perp}(\mathbf{A}) \subseteq \mathbb{Z}^m$, so both lattices are full rank. The duality and determinant properties of these lattices are summarized below.

Lemma 1 ([16, Lemma 2]). *Let* $\mathbf{A}$ *be as in Definition 3. Then*

$$\big(\mathcal{L}_q(\mathbf{A})\big)^{\vee} = \frac{1}{q}\mathcal{L}_q^{\perp}(\mathbf{A}), \qquad \big(\mathcal{L}_q^{\perp}(\mathbf{A})\big)^{\vee} = \frac{1}{q}\mathcal{L}_q(\mathbf{A}).$$

Lemma 2 ([16, Lemma 3]). *Let* $\mathbf{A}$ *be as in Definition 3. Denote by* r *and* r' *the ranks of* $p\mathbf{A}$ *mod* q *and mod* p*, respectively. Then*

$$\det\big(\mathcal{L}_q(\mathbf{A})\big) = p^{-r'}q^{m-r} \geq p^{-k}q^{m-n}.$$

2.4 Periodic Gaussian Function and Fourier Transform

The periodic Gaussian function serves as a tool for comparing candidates, assigning larger values to points near the lattice and smaller values to points farther away. Specifically, let $\rho_s(\mathbf{x}) = \exp\left(-\frac{\pi\|\mathbf{x}\|^2}{s^2}\right)$ be the Gaussian function with parameter $s > 0$, then the periodic Gaussian function of an m-dimensional lattice $\mathcal{L}$ with parameter $s > 0$ is defined by

$$f_{\mathcal{L},s}(\mathbf{t}) \;=\; \frac{\rho_s(\mathbf{t}+\mathcal{L})}{\rho_s(\mathcal{L})} \;\in\; (0,1] \text{ for } \mathbf{t} \in \mathbb{R}^m, \tag{1}$$

where $\mathbf{t}+\mathcal{L} = \{\mathbf{t}+\mathbf{x} : \mathbf{x} \in \mathcal{L}\}$ and $\rho_s(\mathcal{L}) = \sum_{\mathbf{x}\in\mathcal{L}} \rho_s(\mathbf{x})$. The value of $f_{\mathcal{L},s}(\mathbf{t})$ is higher when $\mathbf{t}$ is closer to $\mathcal{L}$, and decreases as the distance to $\mathcal{L}$ increases.

We also denote by $D_{\mathcal{L},s}$ the discrete Gaussian distribution over $\mathcal{L}$ with parameter s, defined as $D_{\mathcal{L},s}(\mathbf{v}) = \frac{\rho_s(\mathbf{v})}{\rho_s(\mathcal{L})}$ for $\mathbf{v} \in \mathcal{L}$. The following approximation lemma is the bridge from the periodic Gaussian function to a computable empirical score, which will later be used to distinguish the correct guess from the wrong ones; see Section 3.1 for details.

Lemma 3 ([1, Lemma 1.2] and [14, Lemma 6]). *Let* $\mathcal{L}$ *be an* m*-dimensional lattice,* $s > 0$*, and let* $f_{\mathcal{L},s}$ *be defined as in* (1)*. Let* N *be an integer,* $\delta > 0$ *and* $X \subseteq \mathbb{R}^m$ *a finite set. Let* $Z = (\mathbf{z}_1, \ldots, \mathbf{z}_N)$ *be a list of vectors in the dual lattice chosen randomly and independently from the distribution* $D_{\mathcal{L}^{\vee},1/s}$*. Then with probability at least* $1 - |X|2^{-\Omega(N\delta^2)}$*,*

$$g_Z(\mathbf{x}) \triangleq \mathrm{Re}\left(\frac{1}{N}\sum_{i=1}^{N}\exp(2\pi\mathbf{i}\langle\mathbf{x},\mathbf{z}_i\rangle)\right) = \frac{1}{N}\sum_{i=1}^{N}\cos(2\pi\langle\mathbf{x},\mathbf{z}_i\rangle)$$

satisfies that $|g_Z(\mathbf{x}) - f_{\mathcal{L},s}(\mathbf{x})| \leq \delta$ *for all* $\mathbf{x} \in \mathcal{L} + X$*.*

We will use the following bound, stated in [16, Corollary 1], to upper-bound $f_{\mathcal{L},s}$ at points far from the lattice.

Lemma 4 ([16, Corollary 1]). *Let $\mathcal{L} \subset \mathbb{R}^m$ be an m-dimensional lattice, $\mathbf{x} \in \mathbb{R}^m$, and set $r_{\mathbf{x}} = \operatorname{dist}(\mathbf{x}, \mathcal{L})$ and $\tau = s\sqrt{\frac{m}{2\pi}}$. If $r_{\mathbf{x}} \geq \tau$, then*

$$\rho_s(r_{\mathbf{x}}) \;\leq\; f_{\mathcal{L},s}(\mathbf{x}) \;\leq\; \left(\tfrac{r_{\mathbf{x}}}{\tau}\right)^m e^{m/2}\, \rho_s(r_{\mathbf{x}}).$$

Moreover, the upper bound $\left(\frac{r_{\mathbf{x}}}{\tau}\right)^m e^{m/2}\rho_s(r_{\mathbf{x}})$ is a decreasing function of $r_{\mathbf{x}}$ over the range $r_{\mathbf{x}} \geq \tau$.

Finally, the computation of the empirical scores over the entire guessing space can be accelerated using the Fast Fourier Transform (FFT). For a function $f : \mathbb{Z}_q^n \to \mathbb{C}$, its discrete Fourier transform is

$$\widehat{f}(\mathbf{v}) = \sum_{\mathbf{u} \in \mathbb{Z}_q^n} f(\mathbf{u}) \exp\left(-\frac{2\pi \mathrm{i}\langle \mathbf{u}, \mathbf{v}\rangle}{q}\right) \text{ for } \mathbf{v} \in \mathbb{Z}_q^n.$$

Using the FFT, all q^n values can be computed simultaneously in time $O(nq^n)$.

2.5 Discrete Gaussian Sampling

Dual attacks require many short dual vectors sampled from a discrete Gaussian distribution. As in [14,16], we also use the MCMC sampler of Wang and Ling [19] to sample from $D_{\mathcal{L}_q^\perp(\mathbf{A}_{\mathrm{dual}}),qs}$. The cost of this sampling step is one of the main terms in the final complexity estimate.

Theorem 1 ([19, Theorem 1, (8), (23) and (24)]). *There is an algorithm that, given a basis $\mathbf{B}$ of a lattice $\mathcal{L} \subset \mathbb{R}^m$, any $\varepsilon > 0$, and any $s > 0$, returns a sample according to some distribution $\mathcal{D}_{\mathcal{L},s,\varepsilon}$ such that $d\big(\mathcal{D}_{\mathcal{L},s,\varepsilon}, \mathcal{D}_{\mathcal{L},s}\big) \leq \varepsilon$. This algorithm runs in time $\ln\left(\frac{1}{\varepsilon}\right) \cdot \frac{1}{\Delta} \cdot \mathrm{poly}(m)$, where*

$$\frac{1}{\Delta} = \frac{\prod_{i=1}^m \rho_{s/\|\widetilde{\mathbf{b}}_i\|}(\mathbb{Z})}{\rho_s(\mathcal{L})},$$

and $\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_m$ are the Gram–Schmidt vectors of $\mathbf{B}$.

In Section 4.3, we instantiate this theorem with $\mathcal{L} = \mathcal{L}_q^\perp(\mathbf{A}_{\mathrm{dual}})$ and width qs, and estimate $1/\Delta$ from the Gram-Schmidt lengths predicted after BKZ reduction.

3 Review of the Provable Dual Attack

In this section, we briefly review two prior works on provable dual attacks. Section 3.1 recalls the core idea of Pouly and Shen [14], while Section 3.2 reviews its modulus switching extension [16] by Qu and Xu. We focus on the core ideas and the overall attack procedure, omitting technical details and proof arguments; we refer the reader to [14,16] for full treatments.

3.1 Provable Dual Attack of [14]

In [14], the secret is partitioned as $\mathbf{s} = (\mathbf{s}_{\text{guess}}^\top, \mathbf{s}_{\text{dual}}^\top)^\top$ with $\mathbf{s}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ and $\mathbf{s}_{\text{dual}} \in \mathbb{Z}_q^{n_{\text{dual}}}$, where $n = n_{\text{guess}} + n_{\text{dual}}$. The matrix is split accordingly into $\mathbf{A} = [\mathbf{A}_{\text{guess}} \mid \mathbf{A}_{\text{dual}}]$, where $\mathbf{A}_{\text{guess}} \in \mathbb{Z}_q^{m \times n_{\text{guess}}}$ and $\mathbf{A}_{\text{dual}} \in \mathbb{Z}_q^{m \times n_{\text{dual}}}$. Then $\mathbf{s}_{\text{guess}}$ is recovered by exhaustively enumerating all candidates with FFT and identifying the correct one. For a candidate value $\widetilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$, observe that

$$\mathbf{b} - \mathbf{A}_{\text{guess}}\widetilde{\mathbf{s}}_{\text{guess}} \;=\; \mathbf{e} + \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \widetilde{\mathbf{s}}_{\text{guess}}) + \mathbf{A}_{\text{dual}}\mathbf{s}_{\text{dual}}. \tag{2}$$

Since $\|\mathbf{e}\|$ is small, two cases are distinguished by the distance of $\mathbf{b} - \mathbf{A}_{\text{guess}}\widetilde{\mathbf{s}}_{\text{guess}}$ to the q-ary lattice $\mathcal{L}_q(\mathbf{A}_{\text{dual}})$.

- If $\widetilde{\mathbf{s}}_{\text{guess}} = \mathbf{s}_{\text{guess}}$, then

$$\operatorname{dist}\big(\mathbf{b} - \mathbf{A}_{\text{guess}}\widetilde{\mathbf{s}}_{\text{guess}},\, \mathcal{L}_q(\mathbf{A}_{\text{dual}})\big) = \operatorname{dist}\big(\mathbf{e},\, \mathcal{L}_q(\mathbf{A}_{\text{dual}})\big) \le \|\mathbf{e}\|,$$

  which is small.
- If $\widetilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$, then under the condition $\lambda_1(\mathcal{L}_q(\mathbf{A})) - \|\mathbf{e}\| > \|\mathbf{e}\|$ established in [14],

$$\begin{aligned}
&\operatorname{dist}\big(\mathbf{b} - \mathbf{A}_{\text{guess}}\widetilde{\mathbf{s}}_{\text{guess}},\, \mathcal{L}_q(\mathbf{A}_{\text{dual}})\big) \\
&= \operatorname{dist}\big(\mathbf{e} + \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \widetilde{\mathbf{s}}_{\text{guess}}), \mathcal{L}_q(\mathbf{A}_{\text{dual}})\big) \\
&= \min\{\|\mathbf{e} + \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \widetilde{\mathbf{s}}_{\text{guess}}) - \mathbf{y}\| : \mathbf{y} \in \mathcal{L}_q(\mathbf{A}_{\text{dual}})\} \\
&\ge \min\{\|\mathbf{e} + \mathbf{x} - \mathbf{y}\| : \mathbf{y} \in \mathcal{L}_q(\mathbf{A}_{\text{dual}}),\ \mathbf{x} \in \mathcal{L}_q(\mathbf{A}_{\text{guess}}) \setminus \mathcal{L}_q(\mathbf{A}_{\text{dual}})\} \\
&\ge \min\{\|\mathbf{x} - \mathbf{y}\| : \mathbf{y} \in \mathcal{L}_q(\mathbf{A}_{\text{dual}}),\ \mathbf{x} \in \mathcal{L}_q(\mathbf{A}_{\text{guess}}) \setminus \mathcal{L}_q(\mathbf{A}_{\text{dual}})\} - \|\mathbf{e}\| \\
&\ge \lambda_1\big(\mathcal{L}_q(\mathbf{A}_{\text{dual}}) + \mathcal{L}_q(\mathbf{A}_{\text{guess}})\big) - \|\mathbf{e}\| = \lambda_1(\mathcal{L}_q(\mathbf{A})) - \|\mathbf{e}\| > \|\mathbf{e}\|,
\end{aligned}$$

which implies that any wrong guess is strictly farther from the lattice than the correct one. To distinguish the two cases, one measures the distance of a vector to $\mathcal{L}_q(\mathbf{A}_{\text{dual}})$ through the periodic Gaussian function $f_{\mathcal{L}_q(\mathbf{A}_{\text{dual}}),\,1/s}$ introduced in Eq. (1):

$$f_{\mathcal{L}_q(\mathbf{A}_{\text{dual}}),\,1/s}\big(\mathbf{b} - \mathbf{A}_{\text{guess}}\widetilde{\mathbf{s}}_{\text{guess}}\big) = \frac{\rho_{1/s}\big(\mathbf{b} - \mathbf{A}_{\text{guess}}\widetilde{\mathbf{s}}_{\text{guess}} + \mathcal{L}_q(\mathbf{A}_{\text{dual}})\big)}{\rho_{1/s}\big(\mathcal{L}_q(\mathbf{A}_{\text{dual}})\big)}. \tag{3}$$

The smaller the distance, the larger the value of (3); hence, the maximum over all $\widetilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ is attained at $\widetilde{\mathbf{s}}_{\text{guess}} = \mathbf{s}_{\text{guess}}$. However, exact computation of Eq. (3) is highly challenging in practice, as it involves an infinite summation. This necessitates an efficient approximation. Let N be a large integer and $Z = (\mathbf{z}_1, \ldots, \mathbf{z}_N)$ be sampled independently from $D_{\mathcal{L}_q^\perp(\mathbf{A}_{\text{dual}}),\,qs}$. Since $\big(\mathcal{L}_q(\mathbf{A}_{\text{dual}})\big)^\vee = \frac{1}{q}\mathcal{L}_q^\perp(\mathbf{A}_{\text{dual}})$, sampling $\mathbf{z} \leftarrow D_{\mathcal{L}_q^\perp(\mathbf{A}_{\text{dual}}),\,qs}$ is equivalent to sampling $\mathbf{z}/q \leftarrow D_{(\mathcal{L}_q(\mathbf{A}_{\text{dual}}))^\vee,\,s}$. Thus, by Lemma 3, for the list Z of independent samples from $D_{\mathcal{L}_q^\perp(\mathbf{A}_{\text{dual}}),\,qs}$, the empirical score

$$g_Z(\mathbf{x}) = \frac{1}{N}\sum_{j=1}^{N} \cos\left(\frac{2\pi\langle \mathbf{z}_j, \mathbf{x}\rangle}{q}\right)$$

provides a highly accurate approximation to Eq. (3) within δ for all candidates simultaneously, with probability at least $1-q^{n_{\text{guess}}}\cdot 2^{-\Omega(N\delta^2)}$. Evaluating all scores simultaneously via the FFT gives a guessing cost of $\text{poly}(m,n)(N+q^{n_{\text{guess}}})$.

3.2 Provable Dual Attack with Modulus Switching of [16]

Qu and Xu [16] replace the original matrix $\mathbf{A}_{\text{guess}}$ with the rounded matrix $\frac{q}{p}\left\lfloor \frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil$ for a much smaller prime p, which reduces the enumeration space from $\mathbb{Z}_q^{n_{\text{guess}}}$ to $\mathbb{Z}_p^{n_{\text{guess}}}$. Accordingly, for a guess $\widetilde{\mathbf{s}}_{\text{guess}}\in\mathbb{Z}_p^{n_{\text{guess}}}$, one has

$$\begin{aligned}\mathbf{b}-\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\widetilde{\mathbf{s}}_{\text{guess}} &= \mathbf{e}+\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil(\mathbf{s}_{\text{guess}}-\widetilde{\mathbf{s}}_{\text{guess}})\\ &\quad+\left(\mathbf{A}_{\text{guess}}-\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\right)\mathbf{s}_{\text{guess}}+\mathbf{A}_{\text{dual}}\mathbf{s}_{\text{dual}}+q\mathbf{u}\end{aligned}$$

for some $\mathbf{u}\in\mathbb{Z}^m$. Rather than treating the rounding error introduced by modulus switching as an additional noise term, the authors of [16] absorb it into a rational q-ary lattice. Concretely, the associated lattice is defined as $\mathcal{L}=\mathcal{L}_q\big([\mathbf{A}_1\mid\mathbf{A}_{\text{dual}}]\big)$, where $\mathbf{A}_1=\left[\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\,\middle|\,\mathbf{A}_{\text{guess}}-\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\right]$. Because of the rounding term, the periodic Gaussian function in (3) is replaced by the following summed variant of the periodic Gaussian function

$$g_{\mathcal{L}_q(\mathbf{A}_{\text{dual}}),1/s}(\mathbf{x})=\sum_{\mathbf{v}\in\mathbb{Z}_q^{n_{\text{guess}}}}\frac{\rho_{1/s}\left(\mathbf{x}-\left(\mathbf{A}_{\text{guess}}-\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\right)\mathbf{v}+\mathcal{L}_q(\mathbf{A}_{\text{dual}})\right)}{\rho_{1/s}(\mathcal{L}_q(\mathbf{A}_{\text{dual}}))}. \tag{4}$$

As before, its empirical score g_Z is computed from N samples $\mathbf{z}_1,\dots,\mathbf{z}_N$ drawn from $D_{\mathcal{L}_q^{\perp}(\mathbf{A}_{\text{dual}}),\,qs}$, namely

$$g_Z(\mathbf{x})=\text{Re}\left(\frac{1}{N}\sum_{j=1}^{N}e^{2\pi\mathrm{i}\langle\mathbf{x},\frac{\mathbf{z}_j}{q}\rangle}\sum_{\mathbf{v}\in\mathbb{Z}_q^{n_{\text{guess}}}}e^{-2\pi\mathrm{i}\left\langle\left(\mathbf{A}_{\text{guess}}-\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\right)\mathbf{v},\frac{\mathbf{z}_j}{q}\right\rangle}\right). \tag{5}$$

The analysis distinguishes two cases. If $\widetilde{\mathbf{s}}_{\text{guess}}\equiv\mathbf{s}_{\text{guess}}\bmod p$, then

$$g_{\mathcal{L}_q(\mathbf{A}_{\text{dual}}),\,1/s}\left(\mathbf{b}-\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\widetilde{\mathbf{s}}_{\text{guess}}\right)\geq\rho_{1/s}(\mathbf{e}).$$

If $\widetilde{\mathbf{s}}_{\text{guess}}\not\equiv\mathbf{s}_{\text{guess}}\bmod p$, then for every $\mathbf{v}\in\mathbb{Z}_q^{n_{\text{guess}}}$, under the condition $\lambda_1(\mathcal{L})\geq\tau+\|\mathbf{e}\|$ with $\tau=\frac{1}{s}\sqrt{\frac{m}{2\pi}}$,

$$\text{dist}\left(\mathbf{b}-\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\widetilde{\mathbf{s}}_{\text{guess}}-\left(\mathbf{A}_{\text{guess}}-\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\right)\mathbf{v},\mathcal{L}_q(\mathbf{A}_{\text{dual}})\right)\geq\lambda_1(\mathcal{L})-\|\mathbf{e}\|\geq\tau.$$

Applying Lemma 4 gives

$$g_{\mathcal{L}_q(\mathbf{A}_{\text{dual}}),\frac{1}{s}}\left(\mathbf{b}-\frac{q}{p}\left\lfloor\frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil\widetilde{\mathbf{s}}_{\text{guess}}\right)\leq q^{n_{\text{guess}}}\left(\frac{\lambda_1(\mathcal{L})-\|\mathbf{e}\|}{\tau}\right)^m e^{m/2}\rho_{1/s}\left(\lambda_1(\mathcal{L})-\|\mathbf{e}\|\right).$$

If, in addition, $\rho_{1/s}(\mathbf{e}) > q^{n_{\text{guess}}} \left(\frac{\lambda_1(\mathcal{L})-\|\mathbf{e}\|}{\tau}\right)^m e^{m/2}\rho_{1/s}(\lambda_1(\mathcal{L}) - \|\mathbf{e}\|)$ holds, then the quantity $g_{\mathcal{L}_q(\mathbf{A}_{\text{dual}}),\,1/s}\left(\mathbf{b} - \frac{q}{p}\left\lfloor \frac{p}{q}\mathbf{A}_{\text{guess}}\right\rceil \widetilde{\mathbf{s}}_{\text{guess}}\right)$ is maximized precisely when $\widetilde{\mathbf{s}}_{\text{guess}} \equiv \mathbf{s}_{\text{guess}} \bmod p$.

Since a single run recovers only $\mathbf{s}_{\text{guess}} \bmod p$, the attack of [16] repeats the subroutine over distinct primes $p_1 < \cdots < p_k < q$, where $\prod_{j=1}^{k} p_j > q$, and recombines the recovered residues via CRT to recover the full secret $\mathbf{s}_{\text{guess}}$. The overall guessing cost is therefore approximately $\text{poly}(m,n)(N + p_k^{n_{\text{guess}}})$, which is dominated by the largest prime p_k.

4 A p-adic Layered Modulus Switching Dual Attack

In this section, we present our p-adic layered modulus switching dual attack, denoted by LaMS. The remainder of this section is organized as follows. In Section 4.1, we introduce the main intuition underlying LaMS. Section 4.2 then formally describes LaMS. Complexity analysis is given in Section 4.3.

4.1 Basic Idea of LaMS

Let $\mathbf{A} = [\mathbf{A}_{\text{guess}} \mid \mathbf{A}_{\text{dual}}] \in \mathbb{Z}_q^{m\times n}$ and $\mathbf{s} = (\mathbf{s}_{\text{guess}}^\top, \mathbf{s}_{\text{dual}}^\top)^\top \in \mathbb{Z}_q^n$, where $n = n_{\text{guess}} + n_{\text{dual}}$. The CRT-based modulus switching strategy recovers the residues of the guessed secret mod several distinct small primes $p_1, \ldots, p_k$. Once $\prod_{j=1}^{k} p_j > q$, the Chinese Remainder Theorem uniquely determines $\mathbf{s}_{\text{guess}} \bmod q$ from the collection of residues $\mathbf{s}_{\text{guess}} \bmod p_j$. This strategy reduces the enumeration space of each individual run from $\mathbb{Z}_q^{n_{\text{guess}}}$ to $\mathbb{Z}_{p_j}^{n_{\text{guess}}}$. However, since the attack has to be carried out for several primes and the cost of the guessing phase grows with the modulus, the overall guessing cost is dominated by the largest CRT prime p_k. To further reduce this cost, we replace the CRT recombination by a layered recovery of p-adic digits with respect to a single fixed small prime p. Instead of recovering $\mathbf{s}_{\text{guess}} \bmod p_j$ for several distinct primes p_j, we first recover the lowest-order residue $\mathbf{s}_{\text{guess}} \bmod p$, subtract its contribution from the LWE samples, and then apply the same modulus switching subroutine to the resulting new target instance. Repeating this procedure reveals the p-adic digits of $\mathbf{s}_{\text{guess}}$ layer by layer, while every layer is searched over the same fixed domain $\mathbb{Z}_p^{n_{\text{guess}}}$.

We represent the guessed secret through its p-adic expansion as

$$\mathbf{s}_{\text{guess}} = \mathbf{s}_{\text{guess}}^{(0)} + p\mathbf{s}_{\text{guess}}^{(1)} + \cdots + p^{\ell-1}\mathbf{s}_{\text{guess}}^{(\ell-1)},$$

where $\mathbf{s}_{\text{guess}}^{(i)} \in \mathbb{Z}_p^{n_{\text{guess}}}$. Fix an arbitrary layer index $i \in \{0, 1, \ldots, \ell-1\}$, where $\ell = \lceil \log_p q \rceil$. Suppose inductively that we have already correctly recovered the lower-order digits $\mathbf{s}_{\text{guess}}^{(0)}, \ldots, \mathbf{s}_{\text{guess}}^{(i-1)}$ (for $i = 0$, the procedure starts by recovering $\mathbf{s}_{\text{guess}}^{(0)}$ directly). We now show how to recover the i-th digit $\mathbf{s}_{\text{guess}}^{(i)}$.

Denote $\mathbf{c}_i = \sum_{j=0}^{i-1} p^j \mathbf{s}_{\text{guess}}^{(j)}$, with $\mathbf{c}_0 = \mathbf{0}$. Then the new target instance, called the i-th layer instance, is given by

$$\mathbf{b}_i \triangleq \mathbf{b} - \mathbf{A}_{\text{guess}}\mathbf{c}_i = \mathbf{B}_i\mathbf{t}_i + \mathbf{A}_{\text{dual}}\mathbf{s}_{\text{dual}} + \mathbf{e} \bmod q, \tag{6}$$

where $\mathbf{t}_i = \sum_{j=i}^{\ell-1} p^{j-i}\mathbf{s}_{\text{guess}}^{(j)}$ and $\mathbf{B}_i = p^i\mathbf{A}_{\text{guess}} \bmod q \in \mathbb{Z}_q^{m\times n_{\text{guess}}}$. Thus, recovering $\mathbf{s}_{\text{guess}}^{(i)}$ is equivalent to recovering $\mathbf{t}_i \bmod p$. Following [16], define $\widehat{\mathbf{B}}_i \triangleq \frac{q}{p}\left\lfloor \frac{p}{q}\mathbf{B}_i \right\rfloor$, $\mathbf{\Delta}_i \triangleq \mathbf{B}_i - \widehat{\mathbf{B}}_i \in \left(\frac{1}{p}\mathbb{Z}/q\mathbb{Z}\right)^{m\times n_{\text{guess}}}$, and let

$$\mathcal{L}_i \triangleq \mathcal{L}_q([\widehat{\mathbf{B}}_i \mid \mathbf{\Delta}_i \mid \mathbf{A}_{\text{dual}}]) \tag{7}$$

denote the i-th layer rational q-ary lattice. We also define the following layer-wise function. For each $i \in \{0,1,\ldots,\ell-1\}$, let

$$G_{L_q(\mathbf{A}_{\text{dual}}),1/s,i}(\mathbf{x}) \triangleq \sum_{\mathbf{v}\in\mathbb{Z}_q^{n_{\text{guess}}}} \frac{\rho_{1/s}(\mathbf{x} - \mathbf{\Delta}_i\mathbf{v} + \mathcal{L}_q(\mathbf{A}_{\text{dual}}))}{\rho_{1/s}(\mathcal{L}_q(\mathbf{A}_{\text{dual}}))}. \tag{8}$$

Given a list $Z = (\mathbf{z}_1,\ldots,\mathbf{z}_N)$ sampled independently from $D_{\mathcal{L}_q^\perp(\mathbf{A}_{\text{dual}}),\,qs}$, define the corresponding empirical score

$$G_{Z,i}(\mathbf{x}) \;=\; \mathrm{Re}\left[\frac{1}{N}\sum_{r=1}^{N}\exp\Big(2\pi\mathrm{i}\,\langle\mathbf{x},\,\tfrac{\mathbf{z}_r}{q}\rangle\Big)\sum_{\mathbf{v}\in\mathbb{Z}_q^{n_{\text{guess}}}}\exp\Big(-2\pi\mathrm{i}\,\langle\mathbf{\Delta}_i\mathbf{v},\,\tfrac{\mathbf{z}_r}{q}\rangle\Big)\right]. \tag{9}$$

A key structural feature distinguishing LaMS from CRT-based recovery is that the modulus p remains fixed across all layers, while the guessing matrix evolves from $\mathbf{A}_{\text{guess}}$ at layer 0 to $p^i\mathbf{A}_{\text{guess}}$ at layer i, and the target vector changes from $\mathbf{b}$ to $\mathbf{b} - \mathbf{A}_{\text{guess}}\mathbf{c}_i$. Moreover, since the dual block $\mathbf{A}_{\text{dual}}$ is unchanged, the same precomputed list of dual samples from $D_{\mathcal{L}_q^\perp(\mathbf{A}_{\text{dual}}),qs}$ can be reused throughout the attack. The difference between the CRT-based strategy and our layered strategy is illustrated in Figure 1.

4.2 Formal Description of LaMS

We now provide the formal description of LaMS. The complete LaMS attack is summarized in Algorithm 1.

At each layer, LaMS invokes the same mod-p recovery subroutine; accordingly, the per-layer digit-recovery subroutine coincides with the modulus switching subroutine of [16]. For this reason, we retain only the essential description and omit the repeated procedural details. The single-layer procedure is presented in Algorithm 2. It corresponds exactly to the procedure described in [16, Algorithm 1], instantiated for the i-th layer. Since the input matrix is $\mathbf{B}_i = p^i\mathbf{A}_{\text{guess}}$ and the target vector is $\mathbf{b}_i = \mathbf{b} - \mathbf{A}_{\text{guess}}\mathbf{c}_i$ with $\mathbf{c}_i = \sum_{j=0}^{i-1} p^j\mathbf{s}_{\text{guess}}^{(j)}$, and since p and q are distinct primes, the scalar p^i is invertible mod q. Hence $\mathbf{B}_i$ is distributed identically to $\mathbf{A}_{\text{guess}}$ mod q and remains independent of $\mathbf{A}_{\text{dual}}$. Consequently, in direct analogy with [16, Theorem 4], we obtain the following single-layer result, stated in Theorem 2.

Theorem 2. *(Adapted from [16, Theorem 4]) Let the notation be defined as above. For a fixed $i \in \{0,1,\ldots,\ell-1\}$, if $[\mathbf{B}_i \mid \mathbf{A}_{\text{dual}}]$ has full rank over $\mathbb{Z}_q$,*

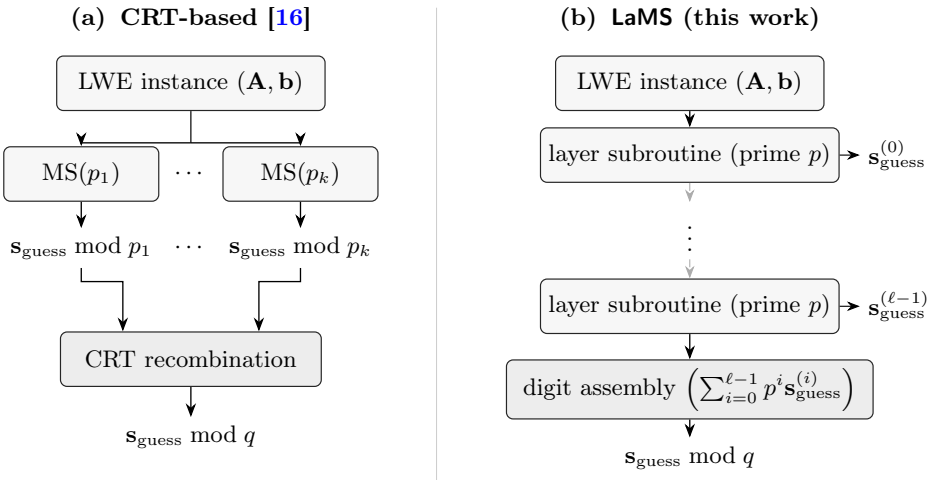


Fig. 1. Comparison of secret-recovery strategies. **(a)** The CRT-based modulus switching dual attack of [16]. **(b)** LaMS fixes a single prime p, recovers the p-adic digits successively, and updates the samples to new target instances between layers.

$(p\mathbf{B}_i) \bmod q$ has full rank over $\mathbb{Z}_p$, $\lambda_1(\mathcal{L}_i) \geq \tau + \|\mathbf{e}\|$, and

$$\rho_{1/s}(\mathbf{e}) - q^{n_{\text{guess}}} \left(\frac{\lambda_1(\mathcal{L}_i) - \|\mathbf{e}\|}{\tau} \right)^m e^{m/2}\, \rho_{1/s}(\lambda_1(\mathcal{L}_i) - \|\mathbf{e}\|) > 2\delta$$

for some $\delta > 0$, where $\mathcal{L}_i$ is defined in Eq. (7) and $\tau = \frac{1}{s}\sqrt{\frac{m}{2\pi}}$, then Algorithm 2 outputs $\mathbf{s}^{(i)}_{\text{guess}}$ with probability at least $1 - p^{n_{\text{guess}}} \cdot 2^{-\Omega(\frac{N\delta^2}{q^{2n_{\text{guess}}}})}$ over the choice of Z.

We are now ready to state the main result for LaMS, given in Theorem 3.

Theorem 3. *Let the notation be defined as above. For all $i \in \{0, 1, \ldots, \ell-1\}$, if $[\mathbf{B}_i \mid \mathbf{A}_{\text{dual}}]$ has full rank over $\mathbb{Z}_q$, $(p\mathbf{B}_i) \bmod q$ has full rank over $\mathbb{Z}_p$ [2], $\lambda_1(\mathcal{L}_i) \geq \tau + \|\mathbf{e}\|$, and*

$$\rho_{1/s}(\mathbf{e}) - q^{n_{\text{guess}}} \left(\tfrac{\lambda_1(\mathcal{L}_i) - \|\mathbf{e}\|}{\tau} \right)^m e^{m/2}\, \rho_{1/s}(\lambda_1(\mathcal{L}_i) - \|\mathbf{e}\|) \;>\; 2\delta$$

for some $\delta > 0$, where $\mathcal{L}_i$ is defined in Eq. (7) and $\tau = \frac{1}{s}\sqrt{\frac{m}{2\pi}}$, then Algorithm 1 outputs $\mathbf{s}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ with probability at least $1 - \ell \cdot p^{n_{\text{guess}}} \cdot 2^{-\Omega(\frac{N\delta^2}{q^{2n_{\text{guess}}}})}$ over the choice of Z.

Proof. Let $\mathbf{c}_\ell = \sum_{i=0}^{\ell-1} p^i \widehat{\mathbf{s}}^{(i)}_{\text{guess}}$ be the output of Algorithm 1. If every layer outputs the correct digit, namely $\widehat{\mathbf{s}}^{(i)}_{\text{guess}} = \mathbf{s}^{(i)}_{\text{guess}}$ for all $i \in \{0, 1, \ldots, \ell-1\}$,

[2] We discuss sufficient conditions under which the required rank conditions of Theorem 2 and 3 hold in Appendix A.

Algorithm 1 LaMS: Layered Modulus Switching

Require: m, $n = n_{\text{guess}} + n_{\text{dual}}$; primes q, p with $p < q$; $N \in \mathbb{N}$; LWE instance $(\mathbf{A}, \mathbf{b})$ with $\mathbf{A} = [\mathbf{A}_{\text{guess}} \mid \mathbf{A}_{\text{dual}}]$; list $Z = \{\mathbf{z}_1, \ldots, \mathbf{z}_N\}$ sampled independently from $D_{\mathcal{L}_q^\perp(\mathbf{A}_{\text{dual}}),\, qs}$ with $s > 0$.

Ensure: The guessed part $\mathbf{s}_{\text{guess}}$ of the secret, mod q.

$\ell \leftarrow \lceil \log_p q \rceil$; $\mathbf{c}_0 \leftarrow \mathbf{0}$.
for $i = 0, 1, \ldots, \ell - 1$ **do**
  $\mathbf{b}_i \leftarrow \mathbf{b} - \mathbf{A}_{\text{guess}}\mathbf{c}_i \bmod q$.
  $\mathbf{B}_i \leftarrow p^i \mathbf{A}_{\text{guess}} \bmod q$.
  $\widehat{\mathbf{s}}_{\text{guess}}^{(i)} \leftarrow$ LaMS-layer applied to $([\mathbf{B}_i \mid \mathbf{A}_{\text{dual}}], \mathbf{b}_i)$ and the list Z with prime p.
  $\mathbf{c}_{i+1} \leftarrow \mathbf{c}_i + p^i \widehat{\mathbf{s}}_{\text{guess}}^{(i)}$.
end for
return $\mathbf{c}_\ell \bmod q$.

Algorithm 2 LaMS-layer: one layer of LaMS

Require: primes q, p with $p < q$; a layer index $i \in \{0, 1, \ldots, \ell - 1\}$; matrix $\mathbf{B}_i \in \mathbb{Z}_q^{m \times n_{\text{guess}}}$ and $\mathbf{A}_{\text{dual}} \in \mathbb{Z}_q^{m \times n_{\text{dual}}}$; a vector $\mathbf{b}_i \in \mathbb{Z}_q^m$; a list $Z = \{\mathbf{z}_1, \ldots, \mathbf{z}_N\}$ sampled independently from $D_{\mathcal{L}_q^\perp(\mathbf{A}_{\text{dual}}), qs}$ with $s > 0$.

Ensure: The i-th p-adic digit $\mathbf{s}_{\text{guess}}^{(i)}$ of $\mathbf{s}_{\text{guess}}$.

Compute $G_{Z,i}\left(\mathbf{b}_i - \frac{q}{p}\lfloor \frac{p}{q}\mathbf{B}_i \rfloor \widetilde{\mathbf{s}}_{\text{guess}}^{(i)}\right)$ for all $\widetilde{\mathbf{s}}_{\text{guess}}^{(i)} \in \mathbb{Z}_p^{n_{\text{guess}}}$ via FFT, where $G_{Z,i}$ is defined in Eq. (9).
return the $\widetilde{\mathbf{s}}_{\text{guess}}^{(i)} \in \mathbb{Z}_p^{n_{\text{guess}}}$ maximising $G_{Z,i}\left(\mathbf{b}_i - \frac{q}{p}\lfloor \frac{p}{q}\mathbf{B}_i \rfloor \widetilde{\mathbf{s}}_{\text{guess}}^{(i)}\right)$.

then $\mathbf{c}_\ell = \mathbf{s}_{\text{guess}} \bmod q$. By Theorem 2, each invocation of LaMS-layer fails with probability at most $p^{n_{\text{guess}}} 2^{-\Omega(\frac{N\delta^2}{q^{2n_{\text{guess}}}})}$. Taking a union bound over the ℓ layers, the probability that any layer fails is at most $\ell p^{n_{\text{guess}}} 2^{-\Omega(\frac{N\delta^2}{q^{2n_{\text{guess}}}})}$. Therefore, Algorithm 1 outputs the correct secret $\mathbf{s}_{\text{guess}}$ with probability at least $1 - \ell p^{n_{\text{guess}}} 2^{-\Omega(\frac{N\delta^2}{q^{2n_{\text{guess}}}})}$. □

Remark 1. In contrast to the setting of [16], where distinct primes are used, our construction employs the same prime p at every layer. This uniformity has the benefit that the rational q-ary lattices arising in all layers share the same determinant. More precisely, under the rank conditions of Theorem 3, following the argument of [16, Lemma 16], it can be deduced that $\det(\mathcal{L}_i) = p^{-n_{\text{guess}}} q^{m-n}$ for all $i \in \{0, 1, \ldots, \ell - 1\}$. Consequently, by the Gaussian heuristic, $\lambda_1(\mathcal{L}_i) \approx \sqrt{\frac{m}{2\pi e}} \cdot \det(\mathcal{L}_i)^{1/m}$ is approximately the same for every $i \in \{0, 1, \ldots, \ell - 1\}$. Therefore, it suffices to use $\lambda_1(\mathcal{L}_0)$ as a representative estimate for all $\lambda_1(\mathcal{L}_i)$.

4.3 Complexity Estimation

In this section, we estimate the cost of LaMS. Our analysis proceeds in three steps: first, we determine the parameters required to ensure success with high

probability; second, we bound the cost of the two main stages of the attack, namely sampling short dual vectors and performing the layered search in Algorithm 1; and third, we instantiate the resulting estimates for Kyber.

Parameter selection. To make the failure probability $\ell p^{n_{\text{guess}}} 2^{-\Omega(\frac{N\delta^2}{q^{2n_{\text{guess}}}})}$ negligible, the parameters δ and N must be chosen appropriately. For the rank conditions that must hold simultaneously across all layers—that is, $[\mathbf{B}_i \mid \mathbf{A}_{\text{dual}}]$ has full rank over $\mathbb{Z}_q$ and $(p\mathbf{B}_i) \bmod q$ has full rank over $\mathbb{Z}_p$ for all $i \in \{0, 1, \ldots, \ell-1\}$—these conditions are simultaneously satisfied with probability at least

$$1 - \frac{2}{(q-1)q^{m-n}} - \ell \cdot \frac{2}{p-1} \cdot \frac{p^{n_{\text{guess}}}(\lfloor q/p \rfloor + 1)^m}{q^m},$$

as established in Appendix A. This probability is exponentially close to 1 when $m > 1.5n$. In particular, for every $i \in \{0, 1, \ldots, \ell-1\}$, we also require that

$$\lambda_1(\mathcal{L}_i) \geq \tau + \|\mathbf{e}\|, \quad \rho_{1/s}(\mathbf{e}) - q^{n_{\text{guess}}} \left(\frac{\lambda_1(\mathcal{L}_i) - \|\mathbf{e}\|}{\tau}\right)^m e^{m/2}\, \rho_{1/s}(\lambda_1(\mathcal{L}_i) - \|\mathbf{e}\|) > 2\delta. \tag{10}$$

We therefore fix these parameters explicitly before turning to the complexity analysis.

Assume that the error vector satisfies $\|\mathbf{e}\| \lesssim \sigma_e \sqrt{\frac{m}{2\pi}}$, which holds with overwhelming probability for an error distribution of width σ_e. Following [16], we choose $\delta = \frac{1}{10}\,\rho_{1/s}\big(\sigma_e\sqrt{\frac{m}{2\pi}}\big) = \frac{1}{10}\,e^{-ms^2\sigma_e^2/2}$. Since the overall failure probability is bounded by $\ell p^{n_{\text{guess}}} 2^{-\Omega(\frac{N\delta^2}{q^{2n_{\text{guess}}}})}$, choosing $N = \frac{(\text{poly}(m)+\log(\ell p^{n_{\text{guess}}}))q^{2n_{\text{guess}}}}{\delta^2}$ makes this bound at most $2^{-\text{poly}(m)}$, and so LaMS succeeds with overwhelming probability.

It remains to ensure that condition (10) holds across all ℓ layers. By Remark 1, the quantities $\lambda_1(\mathcal{L}_i)$ are approximately equal under the Gaussian heuristic, so it suffices to enforce the condition for the representative layer $\mathcal{L}_0$; that is, it suffices to require that $\lambda_1(\mathcal{L}_0) \geq \tau + \|\mathbf{e}\|$ together with

$$\rho_{1/s}(\mathbf{e}) - q^{n_{\text{guess}}} \left(\frac{\lambda_1(\mathcal{L}_0) - \|\mathbf{e}\|}{\tau}\right)^m e^{m/2} \rho_{1/s}(\lambda_1(\mathcal{L}_0) - \|\mathbf{e}\|) > 2\delta. \tag{11}$$

Condition (11) is awkward to handle directly. We therefore invoke the following lemma from [16, Theorem 6], which replaces (11) by a simpler sufficient lower bound on $\lambda_1(\mathcal{L}_0)$.

Lemma 5. *Let* $\tau = \frac{1}{s}\sqrt{\frac{m}{2\pi}}$ *and* $\delta = \frac{1}{10}\rho_{1/s}\big(\sigma_e\sqrt{\frac{m}{2\pi}}\big)$. *Suppose that* $m \geq 500$, $\frac{n_{\text{guess}} \ln q}{m} \leq \frac{4}{5}$, *and* $\|\mathbf{e}\| \leq \sigma_e\sqrt{\frac{m}{2\pi}}$. *Then the condition* $\lambda_1(\mathcal{L}_0) \geq \big(1+\sqrt{1+s^2\sigma_e^2}\big)\tau + \|\mathbf{e}\|$ *implies that*

$$\rho_{1/s}(\mathbf{e}) - q^{n_{\text{guess}}} \left(\frac{\lambda_1(\mathcal{L}_0) - \|\mathbf{e}\|}{\tau}\right)^m e^{m/2}\, \rho_{1/s}(\lambda_1(\mathcal{L}_0) - \|\mathbf{e}\|) > 2\delta.$$

To turn Lemma 5 into an explicit constraint on s, we substitute the determinant value $\det(\mathcal{L}_0) = p^{-n_{\text{guess}}} q^{m-n}$ into the estimate $\lambda_1(\mathcal{L}_0) \approx \sqrt{\frac{m}{2\pi e}}\, \det(\mathcal{L}_0)^{1/m}$, and insert the result into $\lambda_1(\mathcal{L}_0) \geq \left(1 + \sqrt{1 + s^2\sigma_e^2}\right)\tau + \|\mathbf{e}\|$. This yields the effective-regime inequality

$$q^{1-n/m} \geq \left(\frac{1}{s} + \sqrt{\frac{1}{s^2} + \sigma_e^2} + \sigma_e \right) p^{n_{\text{guess}}/m} \sqrt{e}. \tag{12}$$

Cost of sampling short dual vectors. This stage is identical to that of [14,16]: we compute a basis of $\mathcal{L}_q^{\perp}(\mathbf{A}_{\text{dual}})$, apply BKZ with block size β to obtain a reduced basis $\mathbf{B}$, and then run the MCMC discrete Gaussian sampler of Section 2.5 with parameter qs to draw N short dual vectors. The total cost is

$$T_{\text{sample}} = T_{\text{BKZ}}(m, \beta) + N \cdot T_{\text{MCMC}}\left(\mathcal{L}_q^{\perp}(\mathbf{A}_{\text{dual}}), qs\right),$$

where the BKZ cost is evaluated with the lattice estimator of [4] under the cost model of [3,11] and the sieving oracle of [5]. Following the same estimation method as in [14], and ignoring polynomial factors, we estimate the MCMC cost as $T_{\text{MCMC}} = 1/\Delta$. The quantity $1/\Delta$ is estimated exactly as in [14], subject to the constraint that the width s satisfies the effective-regime inequality (12). As this stage carries over unchanged from prior work, we omit further detail and refer the reader to [14,16].

Cost of Algorithm 1. Algorithm 1 invokes LaMS-layer once for each of the $\ell = \lceil \log_p q \rceil$ layers. For a fixed layer i and every candidate $\widetilde{\mathbf{s}}_{\text{guess}}^{(i)} \in \mathbb{Z}_p^{n_{\text{guess}}}$, it evaluates $G_{Z,i}\left(\mathbf{b}_i - \frac{q}{p}\left\lfloor \frac{p}{q}\mathbf{B}_i \right\rfloor \widetilde{\mathbf{s}}_{\text{guess}}^{(i)}\right)$ by FFT. The inner sum of $G_{Z,i}$ over $\mathbb{Z}_q^{n_{\text{guess}}}$ can be evaluated as a finite geometric sum and is therefore evaluated in polynomial time, while the FFT-based maximization over $\mathbb{Z}_p^{n_{\text{guess}}}$ costs $\text{poly}(m,n) \cdot (N + p^{n_{\text{guess}}})$. Summing over all layers, Algorithm 1 runs in time $\text{poly}(m,n) \cdot \ell\,(N + p^{n_{\text{guess}}})$.

Combining the two stages, the total running time is

$$T_{\text{total}} = T_{\text{BKZ}}(m, \beta) + N \cdot T_{\text{MCMC}}\left(\mathcal{L}_q^{\perp}(\mathbf{A}_{\text{dual}}), qs\right) + \text{poly}(m,n) \cdot \ell\,(N + p^{n_{\text{guess}}}).$$

Correction of the error-parameter substitution. In the course of this work, we identified a minor but consequential inconsistency in the complexity estimates of [14,16]. The error $\mathbf{e}$ in the analysis is sampled according to a discrete Gaussian $D_{\mathbb{Z}_q,\sigma_e}$, whose density is proportional to $\rho_{\sigma_e}(x) = \exp(-\pi x^2/\sigma_e^2)$, rather than a normal distribution. Care is needed when matching these two parameterizations: the coefficient of x^2 in the exponent of a normal density $\mathcal{N}(0,\sigma^2)$ is $-1/(2\sigma^2)$, whereas that of the discrete Gaussian is $-\pi/\sigma_e^2$. Matching the two yields $\sigma = \frac{\sigma_e}{\sqrt{2\pi}}$, or equivalently $\sigma_e = \sigma\sqrt{2\pi}$. In contrast, when the Kyber error follows the centered binomial distribution $\mathbf{B}_\eta$, its standard deviation is $\sigma = \sqrt{\eta/2}$. Hence the corresponding Gaussian-width parameter used in the analysis should be $\sigma_e = \sigma\sqrt{2\pi} = \sqrt{\eta\pi}$. However, the estimates in [14,16] directly substitute the standard deviation $\sqrt{\eta/2}$ for σ_e, which amounts to using a smaller error width and therefore underestimates the attack cost.

Table 1. Comparison of complexity estimates for provable dual attacks on Kyber.

Parameter set	Method	Complexity (bits)	m	n_{guess}	β	s
Kyber-512	[14]	238	1023	19	746	0.0900
	[16]	240	1180	1	750	0.0870
	Ours	240	1180	1	748	0.0874
Kyber-768	[14]	347	1519	29	1129	0.1100
	[16]	355	1697	1	1157	0.1035
	Ours	354	1694	1	1156	0.1038
Kyber-1024	[14]	478	2025	39	1602	0.1100
	[16]	491	2260	1	1640	0.1089
	Ours	490	2261	1	1638	0.1091

After applying this correction, the complexity estimates of all three Kyber parameter sets with the methods of [14,16] and with LaMS (taking $p = 2$) are summarised in Table 1 [3].

For the complexity analysis of our method, recall that, in Kyber, each coordinate of the secret $\mathbf{s}$ belongs to $\{0, \pm 1, \pm 2, \pm 3\}$. Therefore, there is no need to enumerate all $\mathbf{v} \in \mathbb{Z}_q^{n_{\text{guess}}}$ in Eq. (9). Since our p-adic method uses $p = 2$, the parity of each coordinate further restricts its candidates: an even residue has possible values in $\{0, 2, 4, 6\}$, whereas an odd residue has possible values in $\{1, 3, 5\}$. Thus, each guessed coordinate has at most four possible candidates. Accordingly, the success probability in Theorem 2 is at least $1 - p^{n_{\text{guess}}} \exp\left(-\frac{N\delta^2}{4^{2n_{\text{guess}}}}\right)$. Hence, it suffices to choose $N = \frac{\left(\text{poly}(m) + \log(\ell p^{n_{\text{guess}}})\right) 4^{2n_{\text{guess}}}}{\delta^2}$. Following [14,16], we omit polynomial factors in the resulting complexity estimates. All complexity estimates are consistently rounded up to the next integer bit, and our estimates and the corrected estimates of [16] follow the same convention.

5 Conclusion

We proposed a new provable dual attack that fixes a single small prime p and recovers the guessed part of the LWE secret digit by digit through its p-adic expansion. All $\ell = \lceil \log_p q \rceil$ layers share the same enumeration space $\mathbb{Z}_p^{n_{\text{guess}}}$ and a single pre-generated list of dual samples, replacing the CRT recombination of [16] and lowering the guessing cost from $O(\text{poly}(m, n)(N + \sum_{j=1}^{k} p_j^{n_{\text{guess}}}))$ in the CRT-based attack to $O(\text{poly}(m, n)\lceil \log_p q \rceil (N + p^{n_{\text{guess}}}))$, where $p \ll p_k$.

[3] The complexity reported for [14] is based on the latest revised version of their work [13]. Note that the complexity of [16] has also been revised in their latest version [15]. However, the error of [15] in the substitution of σ_e remains uncorrected. We therefore correct this issue ourselves, and the resulting complexity is collected in Table 1.

Along the way we identified and corrected a parameter inconsistency in [14,16]. After applying this correction, LaMS reduces the estimated attack cost by up to 1 bit over the corrected [16] across all three Kyber parameter sets. In Appendix B, we further verified that the effective regime of LaMS is disjoint from the contradictory regime of [7], with a gap factor that can be made close to 2, matching the precise characterisation established in [16].

We also highlight two interesting open questions. One is whether modulus switching in provable dual attacks can be replaced by a coding-based decoder, such as the polar-code decoder of [6], thereby reducing the search complexity from $q^{n_{\text{guess}}}$ to $q^{n'}$ for some $n' \ll n_{\text{guess}}$, while retaining a rigorous correctness guarantee. Another open question is whether the factor $q^{2n_{\text{guess}}}$ in the success-probability bound of Theorem 2 can be reduced through a tighter analysis or a better-designed scoring function, potentially yielding a further improvement in the concrete complexity.

Acknowledgments. We thank the anonymous reviewers of Inscrypt 2026 for their constructive feedback. And we sincerely thank the authors of the two preceding works, Amaury Pouly, Yixin Shen and Hongyuan Qu, Guangwu Xu, for the helpful discussions and exchanges regarding their results. We are particularly grateful to Pouly and Shen for generously sharing their code and updated estimates, and for providing careful and detailed responses to our questions.

References

1. Aharonov, D., Regev, O.: Lattice problems in NP ∩ coNP. J. ACM **52**(5), 749–765 (2005). https://doi.org/10.1145/1089023.1089025
2. Albrecht, M.R.: On dual lattice attacks against small-secret LWE and parameter choices in HElib and SEAL. In: Coron, J.S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10211, pp. 103–129. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56614-6_4
3. Albrecht, M.R., Gheorghiu, V., Postlethwaite, E.W., Schanck, J.M.: Estimating quantum speedups for lattice sieves. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020. LNCS, vol. 12492, pp. 583–613. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64834-3_20
4. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. Journal of Mathematical Cryptology **9**(3), 169–203 (2015). https://doi.org/10.1515/jmc-2015-0016
5. Becker, A., Ducas, L., Gama, N., Laarhoven, T.: New directions in nearest neighbor searching with applications to lattice sieving. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 10–24. SODA '16, Society for Industrial and Applied Mathematics, USA (2016). https://doi.org/10.5555/2884435.2884437
6. Carrier, K., Meyer-Hilfiger, C., Shen, Y., Tillich, J.P.: Assessing the impact of a variant of MATZOV's dual attack on Kyber. In: Tauman Kalai, Y., Kamara, S.F. (eds.) CRYPTO 2025. LNCS, vol. 16000, pp. 444–476. Springer, Cham (2025). https://doi.org/10.1007/978-3-032-01855-7_15

7. Ducas, L., Pulles, L.N.: Does the dual-sieve attack on learning with errors even work? In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023. LNCS, vol. 14083, pp. 37–69. Springer Nature Switzerland, Cham (2023). https://doi.org/10.1007/978-3-031-38548-3_2
8. Espitau, T., Joux, A., Kharchenko, N.: On a dual/hybrid approach to small secret LWE. In: Bhargavan, K., Oswald, E., Prabhakaran, M. (eds.) INDOCRYPT 2020. LNCS, vol. 12578, pp. 440–462. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-65277-7_20
9. Guo, Q., Johansson, T.: Faster dual lattice attacks for solving LWE with applications to CRYSTALS. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021. LNCS, vol. 13093, pp. 33–62. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-92068-5_2
10. Lyubashevsky, V., et al.: CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology (2022), https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022
11. MATZOV: Report on the security of LWE: Improved dual lattice attack (2022), https://zenodo.org/record/6412487
12. Nguyen, P.Q.: Hermite's constant and lattice algorithms. In: Nguyen, P.Q., Vallée, B. (eds.) The LLL Algorithm: Survey and Applications, pp. 19–69. Springer, Berlin, Heidelberg (2010). https://doi.org/10.1007/978-3-642-02295-1_2
13. Pouly, A., Shen, Y.: Provable dual attacks on learning with errors. Cryptology ePrint Archive, Paper 2023/1508 (2023), https://eprint.iacr.org/2023/1508
14. Pouly, A., Shen, Y.: Provable dual attacks on learning with errors. In: Joye, M., Leander, G. (eds.) EUROCRYPT 2024. LNCS, vol. 14657, pp. 256–285. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-58754-2_10
15. Qu, H., Xu, G.: On the provable dual attack for LWE by modulus switching. Cryptology ePrint Archive, Paper 2025/859 (2025), https://eprint.iacr.org/2025/859
16. Qu, H., Xu, G.: On the provable dual attack for LWE by modulus switching. In: Hanaoka, G., Yang, B.Y. (eds.) ASIACRYPT 2025. LNCS, vol. 16247, pp. 34–64. Springer, Singapore (2025). https://doi.org/10.1007/978-981-95-5099-9_2
17. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) Proceedings of the 37th Annual ACM Symposium on Theory of Computing. pp. 84–93. ACM Press, Baltimore, MD, USA (2005). https://doi.org/10.1145/1060590.1060603
18. Schwabe, P., et al.: CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology (2022), https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022
19. Wang, Z., Ling, C.: Lattice gaussian sampling by markov chain monte carlo: Bounded distance decoding and trapdoor sampling. IEEE Transactions on Information Theory **65**(6), 3630–3645 (2019). https://doi.org/10.1109/TIT.2019.2901497

Appendix

A Analysis of the Rank Conditions

We first establish the rank conditions used in the single layer. Fix a layer index $i \in \{0, 1, \ldots, \ell - 1\}$. Since $\mathbf{A} = [\mathbf{A}_{\text{guess}} \mid \mathbf{A}_{\text{dual}}] \leftarrow U(\mathbb{Z}_q^{m \times n})$ and $\mathbf{B}_i =$

$p^i\mathbf{A}_{\text{guess}} \bmod q$, and since p and q are distinct primes, p^i is invertible mod q. Hence the map $\mathbf{A}_{\text{guess}} \mapsto p^i\mathbf{A}_{\text{guess}} \bmod q$ is a bijection on $\mathbb{Z}_q^{m\times n_{\text{guess}}}$. It follows that $\mathbf{B}_i$ is uniform over $\mathbb{Z}_q^{m\times n_{\text{guess}}}$ and remains independent of $\mathbf{A}_{\text{dual}}$. Consequently, $[\mathbf{B}_i \mid \mathbf{A}_{\text{dual}}]$ is distributed identically to $\mathbf{A}$, namely $[\mathbf{B}_i \mid \mathbf{A}_{\text{dual}}] \leftarrow U(\mathbb{Z}_q^{m\times n})$. The rank condition for the i-th layer therefore follows; since the remaining argument coincides with that of [15, Appendix A], we omit the details.

Lemma 6 (Rank condition for the i-th layer). *Let $\mathbf{A} = [\mathbf{A}_{\text{guess}} \mid \mathbf{A}_{\text{dual}}] \leftarrow U(\mathbb{Z}_q^{m\times n})$, where $\mathbf{A}_{\text{guess}} \in \mathbb{Z}_q^{m\times n_{\text{guess}}}$ and $\mathbf{A}_{\text{dual}} \in \mathbb{Z}_q^{m\times(n-n_{\text{guess}})}$. For a fixed $i \in \{0,1,\ldots,\ell-1\}$, define $\mathbf{B}_i = p^i\mathbf{A}_{\text{guess}} \bmod q$. Then the probability that $[\mathbf{B}_i \mid \mathbf{A}_{\text{dual}}]$ has full rank over $\mathbb{Z}_q$, and that $(p\mathbf{B}_i) \bmod q$ has full rank over $\mathbb{Z}_p$, is at least*

$$1 - \frac{2}{(q-1)q^{m-n}} - \frac{2}{p-1} \cdot \frac{p^{n_{\text{guess}}}(\lfloor q/p \rfloor + 1)^m}{q^m}.$$

In particular, this probability is exponentially close to 1 if $m > 1.5\,n$.

Then we present the rank conditions that hold simultaneously for all layers. First, observe that the rank condition over $\mathbb{Z}_q$ is common to all layers. Indeed, since p and q are distinct primes, p^i is invertible mod q for every $i \in \{0,\ldots,\ell-1\}$. Hence the transformation $\mathbf{A}_{\text{guess}} \mapsto \mathbf{B}_i = p^i\mathbf{A}_{\text{guess}} \bmod q$ is an invertible linear transformation over $\mathbb{Z}_q$, and therefore the rank of $[\mathbf{B}_i \mid \mathbf{A}_{\text{dual}}]$ over $\mathbb{Z}_q$ is equal to that of $[\mathbf{A}_{\text{guess}} \mid \mathbf{A}_{\text{dual}}]$ over $\mathbb{Z}_q$. Consequently, the full rank condition over $\mathbb{Z}_q$ for all layers reduces to the single condition that $[\mathbf{A}_{\text{guess}} \mid \mathbf{A}_{\text{dual}}]$ has full rank over $\mathbb{Z}_q$, whose failure probability is bounded above by $\frac{2}{(q-1)q^{m-n}}$ (see [15, Appendix A,B] for details).

For each $i \in \{0,1,\ldots,\ell-1\}$, we have $(p\mathbf{B}_i) \bmod q = p^{i+1}\mathbf{A}_{\text{guess}} \bmod q$. Thus, the condition that $(p\mathbf{B}_i) \bmod q$ has full rank over $\mathbb{Z}_p$ is equivalent to requiring $p^{i+1}\mathbf{A}_{\text{guess}} \bmod q$ to have full rank over $\mathbb{Z}_p$. Then taking a union bound over the ℓ layers, the probability that at least one of these $\mathbb{Z}_p$-rank conditions fails is bounded above by $\ell \cdot \frac{2}{p-1} \cdot \frac{p^{n_{\text{guess}}}(\lfloor q/p \rfloor+1)^m}{q^m}$ (see [15, Appendix A,B], for details) [4]. We therefore establish the rank conditions for all layers simultaneously.

Corollary 1 (Rank conditions for all layers). *Let $\mathbf{A} = [\mathbf{A}_{\text{guess}} \mid \mathbf{A}_{\text{dual}}] \leftarrow U(\mathbb{Z}_q^{m\times n})$, where $\mathbf{A}_{\text{guess}} \in \mathbb{Z}_q^{m\times n_{\text{guess}}}$ and $\mathbf{A}_{\text{dual}} \in \mathbb{Z}_q^{m\times(n-n_{\text{guess}})}$. For each $i \in \{0,1,\ldots,\ell-1\}$, define $\mathbf{B}_i = p^i\mathbf{A}_{\text{guess}} \bmod q$. Then the probability that, for all $i \in \{0,1,\ldots,\ell-1\}$, the matrix $[\mathbf{B}_i \mid \mathbf{A}_{\text{dual}}]$ has full rank over $\mathbb{Z}_q$, and $(p\mathbf{B}_i) \bmod q$ has full rank over $\mathbb{Z}_p$, is at least*

$$1 - \frac{2}{(q-1)q^{m-n}} - \ell \cdot \frac{2}{p-1} \cdot \frac{p^{n_{\text{guess}}}(\lfloor q/p \rfloor + 1)^m}{q^m}.$$

In particular, since $\ell = \lceil \log_p q \rceil$ contributes only a polynomial factor, this probability is exponentially close to 1 if $m > 1.5n$.

[4] We clarify that our analysis does not assume independence across layers, and that the union bound applies regardless of dependencies among events.

B Comparison with the Contradictory Regime of [7]

In this section, we start from the effective-regime condition derived in Section 4.3, choose s so as to match [7], and then compare the resulting bound with the contradictory condition. Then we show that the effective regime of LaMS and the contradictory regime are disjoint, with essentially the same constant-factor gap as in [16].

In Section 4.3 we showed that the effective regime of LaMS is

$$q^{1-n/m} \geq \left(\frac{1}{s} + \sqrt{\frac{1}{s^2} + \sigma_e^2} + \sigma_e \right) p^{n_{\text{guess}}/m} \sqrt{e}. \tag{13}$$

Recall that LaMS fixes a single small prime p and reconstructs the secret digit by digit through its p-adic expansion. Consequently, each layer exhausts the same candidate space $\mathbb{Z}_p^{n_{\text{guess}}}$, so the number of guesses per layer is $p^{n_{\text{guess}}}$. The number of layers is $\ell = \lceil \log_p q \rceil$, and $N = \big(\text{poly}(m) + \log(\ell p^{n_{\text{guess}}})\big) q^{2n_{\text{guess}}}/\delta^2$, where $\delta = \frac{1}{10}\, e^{-ms^2\sigma_e^2/2}$.

To match the regime of [7], we choose s such that $\ln T = N\epsilon^2$, where $T = p^{n_{\text{guess}}}$ denotes the number of guesses, $R = \sqrt{4/3}\,\text{GH}(\mathcal{L}_q(\mathbf{A}_{\text{dual}}))$, and $\epsilon = \exp(-\pi R^2 \sigma_e^2)$. Since $\det(\mathcal{L}_q(\mathbf{A}_{\text{dual}})) = q^{m-n_{\text{dual}}}$, the condition $\ln T = N\epsilon^2$ becomes

$$n_{\text{guess}} \ln p = 100\big(\text{poly}(m) + \log(\ell p^{n_{\text{guess}}})\big) q^{2n_{\text{guess}}} \exp\big(ms^2\sigma_e^2 - 2\pi R^2 \sigma_e^2\big)\,.$$

Solving this equation for s yields

$$s = \sqrt{\frac{\ln P}{m\sigma_e^2} + \frac{4}{3e} q^{2-2n_{\text{dual}}/m}}, \tag{14}$$

where $P = \frac{n_{\text{guess}} \ln p}{100\big(\text{poly}(m)+\log(\ell p^{n_{\text{guess}}})\big) q^{2n_{\text{guess}}}}$. Writing $a = 1 - \frac{n_{\text{dual}}}{m} > 0$, Eq. (14) gives

$$s^2 = \frac{4}{3e} q^{2a} + \frac{\ln P}{m\sigma_e^2} = \frac{4}{3e} q^{2a} - \frac{\ln(1/P)}{m\sigma_e^2}.$$

To simplify the expression further, we claim that $\frac{\ln(1/P)}{m\sigma_e^2} \leq \frac{2}{3e} q^{2a}$, which is equivalent to

$$P \geq \exp\left(-\frac{2m\sigma_e^2}{3e} q^{2-\frac{2n_{\text{dual}}}{m}} \right). \tag{15}$$

For our choice of parameters, $m \approx 2n > 2n_{\text{dual}}$. Hence $\frac{2m\sigma_e^2}{3e} \gg 1$ and $q^{2-\frac{2n_{\text{dual}}}{m}} > q$. Therefore, the right-hand side of Eq. (15) is exponentially small, and Eq. (15) is readily satisfied.

It follows that $s^2 \geq \frac{2}{3e} q^{2a}$, and thus $s \geq \sqrt{\frac{2}{3e}}\, q^{1-\frac{n_{\text{dual}}}{m}}$. Since $f(s) = \frac{1}{s} + \sqrt{\frac{1}{s^2} + \sigma_e^2} + \sigma_e$ is decreasing in s, the right-hand side of Eq. (13) is bounded

above by

$$\left(\sqrt{\frac{3e}{2}}\, q^{\frac{n_{\mathrm{dual}}}{m}-1} + \sqrt{\frac{3e}{2} q^{\frac{2n_{\mathrm{dual}}}{m}-2} + \sigma_e^2} + \sigma_e\right) p^{n_{\mathrm{guess}}/m}\sqrt{e}.$$

Since $m \approx 2n > 2n_{\mathrm{dual}}$, we have $\sqrt{\frac{3e}{2}}\, q^{\frac{n_{\mathrm{dual}}}{m}-1} < 1$, and moreover $\frac{3e}{2} q^{\frac{2n_{\mathrm{dual}}}{m}-2} < 1$. Therefore, the right-hand side of Eq. (13) is smaller than

$$(1 + \sqrt{1+\sigma_e^2} + \sigma_e)\, p^{n_{\mathrm{guess}}/m}\sqrt{e}.$$

Using $\sqrt{1+\sigma_e^2} \le 1 + \sigma_e$, together with $p = 2$ and $n_{\mathrm{guess}} \ll m$, we obtain

$$(1 + \sqrt{1+\sigma_e^2} + \sigma_e)\, p^{n_{\mathrm{guess}}/m}\sqrt{e} < (2 + 2\sigma_e)\, p^{n_{\mathrm{guess}}/m}\sqrt{e} < 4(1+\sigma_e),$$

which is smaller than the left-hand side $q^{1-n/m}$ of Eq. (13). This shows that there always exists a choice of s satisfying $\ln T = N\epsilon^2$, under which Eq. (13) holds.

Next, we will compare our effective regime with the contradictory regime identified by Ducas and Pulles [7]. In our notation, this contradiction occurs if there exists at least one index $j \in \{0, 1, \ldots, \ell - 1\}$ such that

$$T_j^{-1/m} \cdot \mathrm{GH}\big(\mathcal{L}_q(\mathbf{A}_{\mathrm{dual}})\big) < \sqrt{\frac{m}{2\pi}} \cdot \sigma_e.$$

In the LaMS setting, we have $T_j = p^{n_{\mathrm{guess}}}$ for all j. Hence the contradictory condition becomes

$$q^{1-\frac{n_{\mathrm{dual}}}{m}} < p^{\,n_{\mathrm{guess}}/m} \cdot \sigma_e\sqrt{e}. \tag{16}$$

Comparing Eq. (13) with Eq. (16) yields

$$q^{1-n_{\mathrm{dual}}/m} > q^{1-n/m} \ge \left(\frac{1}{s} + \sqrt{\frac{1}{s^2} + \sigma_e^2} + \sigma_e\right) p^{n_{\mathrm{guess}}/m}\sqrt{e} > p^{n_{\mathrm{guess}}/m}\sqrt{e}\,\sigma_e.$$

Hence, the two regimes do not overlap. Moreover, since s can be chosen arbitrarily large, the resulting bounds differ only by a factor of two, as in [16].