# What Happens When Integrating Modulus Switching and Lossy Source Coding: A New Dual Attack Variant on LWE

Yechen Li[1] and Qunxiong Zheng[1(✉)]

Information Engineering University, Zhengzhou, 450001, Henan, China
qunxiong_zheng@163.com

**Abstract.** The threat of large-scale quantum computers to classical public-key cryptography has motivated the development of post-quantum cryptographic schemes. Among these, lattice-based constructions have become the mainstream choice in the ongoing NIST standardization process. The security of these schemes typically relies on the hardness of the LWE problem, and the dual-sieve-FFT attack is widely recognized as one of the most effective approaches against it. Recent improvements by MATZOV and Carrier et al. have significantly advanced its efficiency. In this paper, we propose a new variant of the dual-sieve-FFT attack that integrates modulus switching and lossy source coding. We provide a theoretical analysis of the integrated approach and show that the enumeration size in the FFT step can be reduced from $q^{n_{\text{fft}}}$ to $p^{k_{\text{fft}}}$ (with $p < q$, $k_{\text{fft}} < n_{\text{fft}}$), leading to lower FFT cost and decoding cost. When applied to KYBER, our variant achieves better total complexity than the attack using only lossy source coding, although the improvement is modest. More importantly, the decoding and FFT costs are reduced by 1–6 bits and 2–7 bits, respectively, in most parameter settings. These reductions are practically meaningful, especially in multi-target attack scenarios where the FFT and decoding steps are amortized over multiple targets.

**Keywords:** Post-Quantum Cryptography · Learning with Errors · Dual Attack · Modulus Switching · Lossy Source Coding.

## 1 Introduction

### 1.1 Background

The continuous progress in quantum computing makes the advent of large-scale quantum computers increasingly plausible, posing a serious threat to classical public-key cryptography. In response, post-quantum cryptography (PQC) aims at devising cryptographic primitives that remain secure against quantum adversaries. Among the various candidates, lattice-based schemes have emerged as the dominant route in the ongoing standardization efforts. This is owing to their computational efficiency, moderate communication overhead, and rich functionality. In particular, both CRYSTALS-KYBER [14] and CRYSTALS-DILITHIUM [11] are included in the first batch of NIST standards. Their security relies on the hardness of the Module Learning with Errors (MLWE) problem, which is a structured variant of the classical Learning with Errors (LWE) problem.

As the core mathematical foundation of lattice-based cryptography, the LWE problem plays a central role. The precise evaluation of its hardness, and that of its variants, directly determines the security margins of the corresponding cryptographic schemes. Consequently, algorithmic research on solving LWE constitutes a central topic in cryptanalysis. Among the existing attack methods, the dual attack is widely recognized as one of the most effective approaches against LWE.

The dual attack on LWE has undergone a series of refinements over the past decade, evolving from a basic distinguishing method to a sophisticated hybrid framework that combines sieving, FFT, and noise-management techniques. It was first introduced by Micciancio and Regev [13] as a method for distinguishing a bounded distance decoding (BDD) target from a random one by sampling dual vectors. Since LWE can be viewed as a special instance of BDD, the dual attack naturally applies to LWE as well. The core idea is that each sampled dual vector produces a score: BDD targets tend to receive high scores, while random targets receive low ones. To distinguish the two cases reliably, one aggregates the scores over many sampled dual vectors, which in turn requires generating a large number of such vectors. Alkim et al. [3] initially proposed using sieving on the dual lattice to obtain an exponential number of dual vectors, thereby meeting the requirements of the dual attack. Subsequently, Guo et al. [9] combined the dual attack with lattice sieving and the Fast Fourier Transform (FFT), using the FFT to efficiently compute scores for a large number of guesses. This class of attacks is referred to as the dual-sieve-FFT attack [6]. Building on this, MATZOV [12] incorporated modulus switching to reduce the FFT complexity. However, their analysis relies on independence assumptions that have been shown to be problematic [6], and their claimed security reductions for KYBER are not widely accepted. In CRYPTO 2025, Carrier et al. [5] replaced modulus switching with lossy source coding, which not only lowers the FFT cost but also increases the expected score for the correct guess, while avoiding the flawed independence assumptions.

Given the above development, a natural question arises: can modulus switching and lossy source coding be integrated simultaneously into the dual-sieve-FFT framework? (This was already pointed out by Carrier et al. [5] as one of the possible future directions in their concluding remarks.) The two techniques have complementary strengths and weaknesses. Modulus switching introduces no extra computational cost but incurs a rounding error that reduces the score of the correct guess, thereby demanding shorter dual vectors and increasing sampling complexity. Lossy source coding achieves a smaller decoding error, but its decoding step is costly and its FFT efficiency depends on $q$. When $q$ and $n_{\text{fft}}$ are large, modulus switching typically forces $p$ to be very small (e.g., 3 or 5), which severely degrades the score; lossy source coding faces a similar trade-off in $k_{\text{fft}}$—too small leads to large decoding error, too large leaves $q^{k_{\text{fft}}}$ infeasible. Integrating both techniques is therefore worth exploring, as it allows one to choose $p$ and $k_{\text{fft}}$ appropriately to reduce the FFT enumeration size while maintaining a reasonable decoding error, offering more flexibility than either technique alone.

However, the design and efficiency analysis of such an integrated algorithm have now become a critical bottleneck. The challenge is not just to add two techniques together, but to understand whether they can work together in a meaningful way—and this requires a clear analytical framework. The earlier approaches in MATZOV and

Carrier et al. already demand considerable technical sophistication, and extending their approaches to a combined setting seems far from trivial. Moreover, we note that in the results of Carrier et al., the overall complexity is already dominated by the sampling time; it is therefore worth examining whether further improvements to the FFT step can actually have a tangible impact on the total complexity, or whether the sampling bottleneck simply absorbs most of the gains. Recently, the work of [10] provided a unified theoretical lens for dual-sieve-FFT attacks, which offers the necessary analytical tools for approaching this problem. Building on this framework, we are able to carry out the analysis systematically and complete the present work.

### 1.2 Our Contribution

Our contributions are summarized below:

1. **A new dual attack variant on LWE that simultaneously integrates modulus switching and lossy source coding.** We propose a new variant of the dual-sieve-FFT attack that simultaneously integrates modulus switching and lossy source coding. Specifically, we examine two possible orders of composition: modulus-switching-first (MS-LSC) and lossy-source-coding-first (LSC-MS). Our analysis shows that only the MS-LSC route achieves the anticipated performance gain, reducing the enumeration size in the FFT step from $q^{n_{\text{fft}}}$ to $p^{k_{\text{fft}}}$ (with $p < q$, $k_{\text{fft}} < n_{\text{fft}}$). Since the modulus switching is performed first, the generator matrix $\widetilde{\mathbf{G}}$ of the polar code is defined over $\mathbb{Z}_p$, which significantly reduces the decoding time compared with the $\mathbb{Z}_q$-based counterpart. Furthermore, by adopting the unified prediction model of [10], we overcome the analytical challenge of the new variant and extend their model to analyze the expectation of the score under the correct guess—a task that only requires expectation and avoids the problematic independence assumptions. Small-scale experiments confirm that the experimental results are in good agreement with the theoretical predictions, thereby validating the soundness of our analysis.
2. **Flexible parameter selection and improved complexity for KYBER.** Carrier et al. [5] observed that after integrating the two techniques, choosing the intermediate modulus $p$ as a power of two accelerates the FFT. In this work, we target KYBER and consider $p = 512, 1024, 2048$, and $3072$: the first three are powers of two, while $p = 3072$ is chosen for its proximity to $q = 3329$ and its factorization into small primes, both of which benefit FFT efficiency. For these four parameter settings, we provide the optimal complexities for KYBER under various parameter configurations and across the three cost models. In all cases, our results consistently improve upon those of [5], although the improvements are modest. As in [5], the overall time complexity is dominated by the short-vector sampling cost $T_{\text{sample}}$. We emphasize that our approach offers greater flexibility in parameter selection, while achieving total complexity that is comparable to or better than that of [5]. Moreover, although the overall complexity is dominated by the sampling step, the reductions in the per-iteration costs are by no means negligible. In most parameter settings, the FFT time and decoding time are reduced by 1 to 6 bits and 2 to 7 bits, respectively, compared with [5]. These improvements are practically meaningful, especially in scenarios where the sampling cost is amortized over multiple targets: in such settings,

the relative weight of the FFT and decoding steps increases, making their reduction more noticeable. Taken together, these advantages suggest that our variant offers not only competitive total complexity but also tangible practical benefits in real-world cryptanalytic settings.

### 1.3 Organization

The remainder of this paper is organized as follows. Section 2 introduces the necessary background, including lattice background, dual attacks, and the relevant tools used in subsequent sections. Section 3 presents the proposed MS-LSC variant of the dual-sieve-FFT attack, along with a detailed analysis of the two possible integration orders and the score distribution under the correct guess. Section 4 provides a comprehensive complexity analysis of the new variant on KYBER parameters, with detailed parameter tables for different cost models and intermediate moduli. Section 5 concludes the paper and discusses directions for future work. Additional technical details are deferred to Appendix A and B.

## 2 Preliminaries

Let $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{R}$, and $\mathbb{C}$ denote the sets of natural numbers, integers, real numbers, and complex numbers, respectively. For a positive integer $q$, the set of residue classes modulo $q$ is denoted by $\mathbb{Z}_q = \{0, 1, \ldots, q-1\}$. Vectors are denoted by bold lowercase letters and matrices by bold uppercase letters (e.g., $\mathbf{v}$ and $\mathbf{M}$); their transposes are written as $\mathbf{v}^\top$ or $\mathbf{M}^\top$. Unless otherwise specified, all vectors are regarded as column vectors. The concatenation of two row vectors $\mathbf{x}^\top$ and $\mathbf{y}^\top$ is written as $(\mathbf{x}^\top, \mathbf{y}^\top)$, and the inner product of $\mathbf{x}$ and $\mathbf{y}$ is denoted by $\langle \mathbf{x}, \mathbf{y} \rangle$. For a real number $x$, the notation $\lfloor x \rceil$ represents rounding to the nearest integer, with ties (fractional part exactly 0.5) rounded upward. For vectors, rounding is applied componentwise. The rounding error is defined as $\{\mathbf{z}\} \triangleq \mathbf{z} - \lfloor \mathbf{z} \rceil$. Finally, for integers $a \le b$, the set $\{a, a+1, \ldots, b\}$ is denoted by $[a, b]$.

### 2.1 Probability and Distributions

The expectation and variance of a random variable $X$ are denoted by $\mathbb{E}[X]$ and $\mathbb{V}[X] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$, respectively. For a continuous distribution $\mathcal{D}$ with density $f_{\mathcal{D}}$, the expectation of a function $g(X)$ for $X \sim \mathcal{D}$ is given by $\mathbb{E}[g(X)] = \int_{\mathrm{supp}(\mathcal{D})} g(x) f_{\mathcal{D}}(x)\, \mathrm{d}x$, whereas for a discrete distribution with probability mass function $p_{\mathcal{D}}$, $\mathbb{E}[g(X)] = \sum_{x \in \mathrm{supp}(\mathcal{D})} g(x) p_{\mathcal{D}}(x)$, where $\mathrm{supp}(\mathcal{D})$ denotes the support of the distribution, i.e. the set of points where the density or mass function is non-zero. This representation will be used in later sections to compute expectations.

In particular, we denote by $\mathcal{N}(\mu, \sigma^2)$ the normal distribution with mean $\mu$ and variance $\sigma^2$, and by $\mathcal{U}(\mathcal{X})$ the uniform distribution over a bounded set $\mathcal{X}$. Clearly, both $\mathcal{N}(0, \sigma^2)$ and $\mathcal{U}([-a, a))$ are symmetric about zero for any $\sigma, a > 0$. Another symmetric distribution that plays a central role in lattice-based cryptography is the centered binomial distribution, which we define next.

**Definition 1 (Centered Binomial Distribution).** *The centered binomial distribution $\mathcal{B}_\eta$ of parameter $\eta \in \mathbb{Z}$ is defined as $\mathcal{B}_\eta \sim \sum_{i=1}^{\eta}(X_i - Y_i)$, where $X_i$ and $Y_i$ are independent and identically distributed (i.i.d.) uniform random variables over $\{0, 1\}$. In particular, for random $X \sim \mathcal{B}_\eta$, $\Pr(X = j) = \frac{1}{2^{2\eta}}\binom{2\eta}{\eta+j}$ for all $j \in [-\eta, \eta]$.*

It is clear that $\mathcal{B}_\eta$ is a symmetric distribution with a mean 0 and standard deviation $\sigma = \sqrt{\eta/2}$.

Expectations of cosine functions of random variables will be used frequently in our analysis. By the integral formula for expectation, the cosine addition formula, and the fact that $\mathbb{E}[\sin(X)] = 0$ for any symmetric distribution, we can obtain the following lemma.

**Lemma 1.** *If $X \sim \mathcal{U}([-\frac{1}{2}, \frac{1}{2}))$, then for any constant $t \in \mathbb{R}$, $\mathbb{E}[\cos(tX)] = 2\sin(t/2)/t$. Moreover, if $X$ and $Y$ are independent random variables, both following symmetric distributions, then $\mathbb{E}[\cos(X + Y)] = \mathbb{E}[\cos(X)] \cdot \mathbb{E}[\cos(Y)]$.*

### 2.2 Lattice Background

Given a full column-rank matrix $\mathbf{B} \in \mathbb{R}^{m\times n}$, the lattice generated by $\mathbf{B}$, denoted by $\Lambda(\mathbf{B})$, is the set of all integer linear combinations of the column vectors of $\mathbf{B}$, namely, $\Lambda(\mathbf{B}) = \{\mathbf{Bx} \mid \mathbf{x} \in \mathbb{Z}^n\}$. The column vectors $\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n$ of $\mathbf{B}$ are called a *basis* of the lattice $\Lambda(\mathbf{B})$, and the integer $n$ is called its rank. For simplicity, we refer to $\mathbf{B}$ itself as the basis of $\Lambda(\mathbf{B})$. We now give the definition that is central to the lattice-based problem considered in this paper.

**Definition 2 (LWE Instance).** *Let $n$ and $m$ be positive integers, $\chi_e$ a probability distribution on $\mathbb{Z}_q$ and $\mathbf{s} \in \mathbb{Z}_q^n$ a secret vector. A pair*

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

*is called an LWE instance, where $\mathbf{A}$ is chosen uniformly at random and $\mathbf{e} \sim \chi_e^m$.*

In modern LWE schemes, both secret and error vectors are short and follow the same distribution $\chi_s = \chi_e$ (zero mean, small $\sigma_e$).

### 2.3 Dual Attack

For an LWE instance $(\mathbf{A}, \mathbf{b} = \mathbf{As} + \mathbf{e} \bmod q)$, the dual-sieve-FFT attack first partitions the matrix $\mathbf{A}$ into three blocks

$$\mathbf{A} = [\mathbf{A}_{\text{lat}} \mid \mathbf{A}_{\text{enu}} \mid \mathbf{A}_{\text{fft}}] \in \mathbb{Z}_q^{m\times n_{\text{lat}}} \times \mathbb{Z}_q^{m\times n_{\text{enu}}} \times \mathbb{Z}_q^{m\times n_{\text{fft}}},$$

where $n = n_{\text{lat}} + n_{\text{enu}} + n_{\text{fft}}$, and correspondingly splits the secret vector as $\mathbf{s} = (\mathbf{s}_{\text{lat}}^\top, \mathbf{s}_{\text{enu}}^\top, \mathbf{s}_{\text{fft}}^\top)^\top \in \mathbb{Z}_q^{n_{\text{lat}}} \times \mathbb{Z}_q^{n_{\text{enu}}} \times \mathbb{Z}_q^{n_{\text{fft}}}$. Consider the lattice basis

$$\mathbf{B} = \begin{bmatrix} \mathbf{I}_m & \mathbf{0} \\ \mathbf{A}_{\text{lat}}^\top & q\mathbf{I}_{n_{\text{lat}}} \end{bmatrix}, \tag{1}$$

applying BKZ reduction and sieving to $\mathbf{B}$ yields a list of short vectors $\mathcal{W} = \{\mathbf{w}_i\}_{i=1}^{N}$, each $\mathbf{w}_i = (\mathbf{x}_i^\top, \mathbf{y}_i^\top)^\top \in \mathbb{Z}_q^m \times \mathbb{Z}_q^{n_{\text{lat}}}$ satisfying $\mathbf{A}_{\text{lat}}^\top \mathbf{x}_i \equiv \mathbf{y}_i \bmod q$. Since both $\mathbf{s}_{\text{lat}}$ and $\mathbf{e}$ are short vectors, it follows that

$$\langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}}\mathbf{s}_{\text{enu}} - \mathbf{A}_{\text{fft}}\mathbf{s}_{\text{fft}} \rangle = \langle \mathbf{x}_i, \mathbf{A}_{\text{lat}}\mathbf{s}_{\text{lat}} + \mathbf{e} \rangle = \langle \mathbf{w}_i, \mathbf{e}' \rangle \bmod q$$

is a small quantity, where $\mathbf{e}' = (\mathbf{e}^\top, \mathbf{s}_{\text{lat}}^\top)^\top$. Conversely, if the guess $(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}})$ differs from the true value, then $\langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}}\widetilde{\mathbf{s}_{\text{enu}}} - \mathbf{A}_{\text{fft}}\widetilde{\mathbf{s}_{\text{fft}}} \rangle \bmod q$ is approximately uniformly distributed over $\mathbb{Z}_q$. In the dual attack, the score is evaluated by enumerating $\widetilde{\mathbf{s}_{\text{enu}}}$ and, for each such guess, traversing all possible $\widetilde{\mathbf{s}_{\text{fft}}}$ candidates via FFT; this choice of parameters is made to balance the cost of the search with that of the short-vector sampling step. For a candidate guess $(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}})$, the score function is defined as

$$F_{\mathcal{W}}(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) = \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left( \frac{2\pi}{q} \left( \langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}}\widetilde{\mathbf{s}_{\text{enu}}} \rangle - \langle \mathbf{A}_{\text{fft}}^\top \mathbf{x}_i, \widetilde{\mathbf{s}_{\text{fft}}} \rangle \right) \right). \tag{2}$$

When the guess is correct, i.e. $(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) = (\mathbf{s}_{\text{enu}}, \mathbf{s}_{\text{fft}})$, each term of $F_{\mathcal{W}}$ is close to 1 and its expectation is significantly positive; in contrast, for an incorrect guess, the terms behave as random values in $(-1, 1)$ and cancel out, giving an expectation near 0. Hence, the score reliably indicates correctness, allowing recovery of $\mathbf{s}_{\text{enu}}$. MATZOV [12] introduced modulus switching, replacing the modulus $q$ with a small prime $p$, so that the enumeration size in the FFT step is reduced from $q^{n_{\text{fft}}}$ to $p^{n_{\text{fft}}}$. The corresponding score function becomes

$$F_{\mathcal{W}}^{(\text{ms})}(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) = \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left( \frac{2\pi}{p} \left( \frac{p}{q} \langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}}\widetilde{\mathbf{s}_{\text{enu}}} \rangle - \langle \lfloor \frac{p}{q} \mathbf{A}_{\text{fft}}^\top \mathbf{x}_i \rceil, \widetilde{\mathbf{s}_{\text{fft}}} \rangle \right) \right). \tag{3}$$

However, this reduction comes at a cost: the rounding operation introduces an error that reduces the score corresponding to the correct guess, thereby lowering the distinguishability between correct and incorrect candidates.

For $\mathbf{A}_{\text{fft}}^\top \mathbf{x}_i$, Carrier et al. [5] instead decode in a linear code $C(\mathbf{G})$ (with $\mathbf{G} \in \mathbb{Z}_q^{n_{\text{fft}} \times k_{\text{fft}}}$) to obtain the nearest codeword $\mathbf{G}\mathbf{u}_i$, and consider

$$\begin{aligned} F_{\mathcal{W}}^{(\text{lsc})}(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left( \frac{2\pi}{q} \left( \langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}}\widetilde{\mathbf{s}_{\text{enu}}} \rangle - \langle \mathbf{G}\mathbf{u}_i, \widetilde{\mathbf{s}_{\text{fft}}} \rangle \right) \right) \\ &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left( \frac{2\pi}{q} \left( \langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}}\widetilde{\mathbf{s}_{\text{enu}}} \rangle - \langle \mathbf{u}_i, \mathbf{G}^\top \widetilde{\mathbf{s}_{\text{fft}}} \rangle \right) \right). \end{aligned} \tag{4}$$

Since $\mathbf{G}^\top \widetilde{\mathbf{s}_{\text{fft}}} \in \mathbb{Z}_q^{k_{\text{fft}}}$, the FFT enumeration reduces to $q^{k_{\text{fft}}}$ candidates. Compared with modulus switching, this approach introduces a smaller decoding error and thus better preserves the score of the correct guess, but at the cost of a non-negligible decoding time in practice.

### 2.4 Fast Fourier Transform

A crucial step in the dual-sieve-FFT attack is the computation of discrete Fourier transforms, whose definition is recalled below.

**Definition 3 (Discrete Fourier Transform).** *The discrete Fourier transform* $\widehat{f}$ *of a function* $f : \mathbb{Z}_q^n \to \mathbb{C}$ *is defined as*

$$\widehat{f}(\mathbf{x}) = \sum_{\mathbf{a} \in \mathbb{Z}_q^n} f(\mathbf{a}) e^{-\frac{2\pi i}{q}\langle \mathbf{x}, \mathbf{a} \rangle}.$$

In practice, multi-dimensional discrete Fourier transforms are computed via the FFT, which reduces the computation to a sequence of one-dimensional transforms. Carrier et al. [5] provide the following overall complexity estimate for the FFT, which is expressed in terms of the cost of a single one-dimensional transform over $\mathbb{Z}_q$.

**Proposition 1 (Cost of FFT).** *The total cost of an FFT over* $\mathbb{Z}_q^n$ *is*

$$C_{\mathrm{FFT}} = C_{\mathrm{add}} \cdot n \cdot q^{n-1} \cdot N_{\mathrm{FFT}}^{(\mathrm{add})}(q) + C_{\mathrm{mul}} \cdot n \cdot q^{n-1} \cdot N_{\mathrm{FFT}}^{(\mathrm{mul})}(q),$$

*where* $N_{\mathrm{FFT}}^{(\mathrm{add})}(q)$ *and* $N_{\mathrm{FFT}}^{(\mathrm{mul})}(q)$ *are the numbers of additions and multiplications required by a one-dimensional FFT over* $\mathbb{Z}_q$*, and* $C_{\mathrm{add}}$ *and* $C_{\mathrm{mul}}$ *denote the per-operation costs of addition and multiplication, respectively.*

Following the convention of [5], we take $C_{\mathrm{add}} = 160$ and $C_{\mathrm{mul}} = 1024$. Moreover, for any given integer $q$, the values of $N_{\mathrm{FFT}}^{(\mathrm{add})}(q)$ and $N_{\mathrm{FFT}}^{(\mathrm{mul})}(q)$ can be obtained using the FFTW software.[1]

### 2.5 Polar Codes

We denote by $[n, k]_p$ a linear code of length $n$ and rank $k$ over $\mathbb{Z}_p$. The polar codes used in our work are of this type. For a polar code of length $n = 2^\ell$ and rank $k$, its generator matrix is constructed as

$$\mathbf{G} \triangleq (\mathbf{K}_1 \otimes \cdots \otimes \mathbf{K}_\ell) \cdot \mathbf{F},$$

where $\otimes$ denotes the Kronecker product, each

$$\mathbf{K}_i = \begin{bmatrix} 1 & 1 \\ \alpha_i & 0 \end{bmatrix}$$

with $\alpha_i \in \mathbb{Z}_p^\times$, i.e. $\gcd(\alpha_i, p) = 1$, chosen uniformly at random, and $\mathbf{F} \in \mathbb{Z}_p^{n \times k}$ is an expansion matrix. Specifically, for any $\mathbf{m} \in \mathbb{Z}_p^k$, the matrix $\mathbf{F}$ maps $\mathbf{m}$ to a vector $\mathbf{Fm} \in \mathbb{Z}_p^n$ by keeping $k$ specified components identical to those of $\mathbf{m}$ and setting the remaining $n - k$ components to zero (these are the frozen positions). If the desired code length is not a power of two, the above construction is punctured by removing the necessary number of rows from $\mathbf{G}$.

In our setting, we take $n = n_{\mathrm{fft}}$ and $k = k_{\mathrm{fft}}$. The construction and the corresponding estimates from [5] are adopted. The decoding distance in this setting is then given by

$$d \approx p^{1 - \frac{k_{\mathrm{fft}}}{n_{\mathrm{fft}}}} \sqrt{\frac{n_{\mathrm{fft}}}{2\pi e}}, \tag{5}$$

[1] The software is available at https://www.fftw.org/. Its usage follows the same procedure as the open-source code of [5]; the numerical values may vary slightly across different machines.

which serves as the lattice analogue of the Gilbert–Varshamov bound.

The following lemma, taken from [5], gives the decoding complexity of such polar codes when using the probabilistic successive cancellation (SC) decoder employed in their work.

**Lemma 2 ([5], Lemma 4.2).** *Let $N_{\mathrm{FFT}}^{(\mathrm{add})}$, $N_{\mathrm{FFT}}^{(\mathrm{mul})}$, $C_{\mathrm{add}}$ and $C_{\mathrm{mul}}$ be defined as in Proposition 1. Then list decoding an $[n_{\mathrm{fft}}, k_{\mathrm{fft}}]_p$ polar code over $\mathbb{Z}_p$ using $L$ probabilistic SC decoders (typically $L = 1$) has time complexity*

$$T_{\mathrm{dec}}(p, n_{\mathrm{fft}}, k_{\mathrm{fft}}) = 3 \cdot L \cdot \left(C_{\mathrm{mul}} \cdot N_{\mathrm{FFT}}^{(\mathrm{mul})}(p) + C_{\mathrm{add}} \cdot N_{\mathrm{FFT}}^{(\mathrm{add})}(p)\right) \cdot n'_{\mathrm{fft}} \log_2(n'_{\mathrm{fft}}),$$

*where $n'_{\mathrm{fft}}$ is the smallest power of 2 greater than $n_{\mathrm{fft}}$.*

## 3 A New Dual Attack Variant on LWE

In this section, we introduce a variant of the dual-sieve-FFT attack that simultaneously integrates modulus switching and lossy source coding, aiming to improve the overall time complexity. We first examine two possible orders for integrating the two techniques and show analytically that only one of them yields a further reduction in the FFT complexity. Based on this order, we then propose a new attack variant. Finally, following [10], we provide a theoretical analysis of the new variant, estimating the expected score corresponding to the correct guess.

### 3.1 Two Routes for integrating the Techniques

To gain insight into the behaviour of the score functions, we first derive their explicit forms when the guess is correct, i.e. when $(\widetilde{\mathbf{s}_{\mathrm{enu}}}, \widetilde{\mathbf{s}_{\mathrm{fft}}}) = (\mathbf{s}_{\mathrm{enu}}, \mathbf{s}_{\mathrm{fft}})$. Under this condition, the score functions from Eqs. (3) and (4) reduce to the following expressions: for the modulus-switching variant,

$$\begin{aligned} F_{\mathcal{W}}^{(\mathrm{ms})}(\mathbf{s}_{\mathrm{enu}}, \mathbf{s}_{\mathrm{fft}}) &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{p}\left(\frac{p}{q}\langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\mathrm{enu}}\mathbf{s}_{\mathrm{enu}}\rangle + \langle\left\{\frac{p}{q}\mathbf{A}_{\mathrm{fft}}^{\top}\mathbf{x}_i\right\} - \frac{p}{q}\mathbf{A}_{\mathrm{fft}}^{\top}\mathbf{x}_i, \mathbf{s}_{\mathrm{fft}}\rangle\right)\right) \\ &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{q}\langle \mathbf{w}_i, \mathbf{e}'\rangle + \frac{2\pi}{p}\left\langle\left\{\frac{p}{q}\mathbf{A}_{\mathrm{fft}}^{\top}\mathbf{x}_i\right\}, \mathbf{s}_{\mathrm{fft}}\right\rangle\right); \end{aligned}$$

for the lossy-source-coding variant,

$$\begin{aligned} F_{\mathcal{W}}^{(\mathrm{lsc})}(\mathbf{s}_{\mathrm{enu}}, \mathbf{s}_{\mathrm{fft}}) &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{q}\left(\langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\mathrm{enu}}\mathbf{s}_{\mathrm{enu}}\rangle + \langle \mathbf{e}_i - \mathbf{A}_{\mathrm{fft}}^{\top}\mathbf{x}_i, \mathbf{s}_{\mathrm{fft}}\rangle\right)\right) \\ &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{q}\left(\langle \mathbf{w}_i, \mathbf{e}'\rangle + \langle \mathbf{e}_i, \mathbf{s}_{\mathrm{fft}}\rangle\right)\right), \end{aligned}$$

where $\mathbf{e}' = (\mathbf{e}^{\top}, \mathbf{s}_{\mathrm{lat}}^{\top})^{\top}$ and $\mathbf{e}_i = \mathbf{A}_{\mathrm{fft}}^{\top}\mathbf{x}_i - \mathbf{G}\mathbf{u}_i$. In both cases, the cosine argument is small for the correct guess, making the expected score significantly larger than zero, and thus enabling one to distinguish the correct guess from incorrect ones.

We now consider integrating modulus switching with lossy source coding. Two natural orders arise: modulus switching first, followed by lossy source coding (denoted MS-LSC), or the reverse, lossy source coding first followed by modulus switching (denoted LSC-MS). We examine both routes by inspecting the corresponding score functions.

**MS-LSC Route.** If modulus switching is applied first, the score function becomes

$$F^1_{\mathcal{W}}(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) = \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{q}\langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}}\widetilde{\mathbf{s}_{\text{enu}}}\rangle - \frac{2\pi}{p}\left\langle \left\lfloor \frac{p}{q}\mathbf{A}^\top_{\text{fft}}\mathbf{x}_i \right\rceil, \widetilde{\mathbf{s}_{\text{fft}}} \right\rangle\right).$$

At this point, the FFT enumeration size is $p^{n_{\text{fft}}}$. To reduce it to $p^{k_{\text{fft}}}$, we apply lossy source coding to decode $\left\lfloor \frac{p}{q}\mathbf{A}^\top_{\text{fft}}\mathbf{x}_i \right\rceil$. Let $\widetilde{\mathbf{G}} \in \mathbb{Z}_p^{n_{\text{fft}} \times k_{\text{fft}}}$ be the generator matrix of the corresponding polar code, and let $\widetilde{\mathbf{u}_i}$ be the information vector of the nearest codeword in $C(\widetilde{\mathbf{G}})$ to $\left\lfloor \frac{p}{q}\mathbf{A}^\top_{\text{fft}}\mathbf{x}_i \right\rceil$. The score function then becomes

$$F^2_{\mathcal{W}}(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) = \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{q}\langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}}\widetilde{\mathbf{s}_{\text{enu}}}\rangle - \frac{2\pi}{p}\langle \widetilde{\mathbf{G}}\widetilde{\mathbf{u}_i}, \widetilde{\mathbf{s}_{\text{fft}}}\rangle\right).$$

To understand how this score function behaves for the correct guess, we rewrite it in a more transparent form. If $(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) = (\mathbf{s}_{\text{enu}}, \mathbf{s}_{\text{fft}})$, by a similar derivation, the score simplifies to

$$F^2_{\mathcal{W}}(\mathbf{s}_{\text{enu}}, \mathbf{s}_{\text{fft}}) = \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{q}\langle \mathbf{w}_i, \mathbf{e}'\rangle + \frac{2\pi}{p}\left\langle \left\{\frac{p}{q}\mathbf{A}^\top_{\text{fft}}\mathbf{x}_i\right\}, \mathbf{s}_{\text{fft}}\right\rangle + \frac{2\pi}{p}\langle \widetilde{\mathbf{e}_i}, \mathbf{s}_{\text{fft}}\rangle\right),$$

where $\widetilde{\mathbf{e}_i} = \left\lfloor \frac{p}{q}\mathbf{A}^\top_{\text{fft}}\mathbf{x}_i \right\rceil - \widetilde{\mathbf{G}}\widetilde{\mathbf{u}_i}$. This expression reveals that the cosine argument consists of three small terms: the first comes from the original LWE error, the second is the rounding error from modulus switching, and the third is the coding error from lossy source coding. Since all three are small, the score under the correct guess is biased towards 1, producing a large expectation that distinguishes it from incorrect guesses. Furthermore, because $\widetilde{\mathbf{G}}$ is defined over $\mathbb{Z}_p$, the decoding step is more efficient than its $\mathbb{Z}_q$-based counterpart, and $\|\widetilde{\mathbf{e}_i}\|$ is shorter. These observations confirm the viability of the MS-LSC route, which we adopt for integrating the two techniques in our new variant.

**LSC-MS Route.** For this route, we only state the final result; the derivation is deferred to Appendix A. Under the correct guess, applying lossy source coding first and then modulus switching yields

$$F^2_{\mathcal{W}}(\mathbf{s}_{\text{enu}}, \mathbf{s}_{\text{fft}}) = \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{q}\left(\langle \mathbf{w}_i, \mathbf{e}'\rangle + \langle \mathbf{e}_i, \mathbf{s}_{\text{fft}}\rangle\right) + \frac{2\pi}{p}\left\langle \left\{\frac{p}{q}\mathbf{u}_i\right\}, \mathbf{G}^\top \mathbf{s}_{\text{fft}}\right\rangle\right),$$

where $\mathbf{e}_i = \mathbf{A}_{\text{fft}}^\top \mathbf{x}_i - \mathbf{G}\mathbf{u}_i$. Since $\mathbf{G}^\top \mathbf{s}_{\text{fft}}$ is not necessarily short, the last term is not guaranteed to be small. This significantly lowers the expectation of the score for the correct guess and weakens the distinguishability. Hence the LSC-MS route is not viable.

Based on the above analysis, we adopt the MS-LSC route and define the score function for our new variant as

$$F_{\mathcal{W}}^{(\text{new})}(\widetilde{\mathbf{s}_{\text{enu}}}, \mathbf{z}) = \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{q}\langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}}\widetilde{\mathbf{s}_{\text{enu}}}\rangle - \frac{2\pi}{p}\langle \widetilde{\mathbf{u}_i}, \mathbf{z}\rangle\right). \tag{6}$$

To facilitate a direct comparison with the results of [5], we adopt the same algorithmic structure as their Algorithm 3.1. To make the algorithm more explicit, and more importantly, to facilitate the subsequent complexity analysis, we briefly describe each subroutine of Algorithm 1 below.

**Algorithm 1:** The MS-LSC based dual attack to solve LWE

**Input** : an LWE instance $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$.
**Parameters:** some positive integers $R, T, n_{\text{enu}}, n_{\text{fft}}, k_{\text{fft}}, n_{\text{lat}}$ and an $[n_{\text{fft}}, k_{\text{fft}}]_p$ linear code with generator matrix $\widetilde{\mathbf{G}}$.
**Output** : the secret vector $\mathbf{s}$.

choose $I_{\text{lat}} \subseteq [1, n]$ with $|I_{\text{lat}}| = n_{\text{lat}}$;
$\mathcal{W} \leftarrow$ `Short_Vectors_Sample`$(\mathbf{A}, I_{\text{lat}})$;
**repeat**
  choose a partition $I_{\text{enu}}$ and $I_{\text{fft}}$ of $[1, n] \setminus I_{\text{lat}}$ with $|I_{\text{enu}}| = n_{\text{enu}}, |I_{\text{fft}}| = n_{\text{fft}}$;
  $\mathbf{A}_{\text{enu}}, \mathbf{A}_{\text{fft}} \leftarrow$ select the columns of $\mathbf{A}$ indexed by $I_{\text{enu}}$ and $I_{\text{fft}}$ respectively;
  $\mathscr{L} \leftarrow$ `Modulus_Switch_with_List_Decode`$(\mathcal{W}, \mathbf{A}_{\text{fft}}, \widetilde{\mathbf{G}})$;
  $V \leftarrow$ `Solve_LWE_with_FFT`$(\mathscr{L})$;
  **if** $V \geq T$ **then**
    $\mathbf{s}_{\text{enu}} \leftarrow \mathbf{0}$;
    $(\mathbf{s}_{\text{fft}}, \mathbf{s}_{\text{lat}}) \leftarrow$ `Sub_LWE_Solver`$([\mathbf{A}_{\text{fft}} \mid \mathbf{A}_{\text{lat}}], \mathbf{b})$;
    **return** $(\mathbf{s}_{\text{fft}}, \mathbf{s}_{\text{lat}}, I_{\text{enu}}, I_{\text{fft}}, I_{\text{lat}})$;
  **end**
**until** *R times*;

**`Short_Vectors_Sample`.** In this step, the columns of $\mathbf{A}$ indexed by $I_{\text{lat}}$ are selected to form the matrix $\mathbf{A}_{\text{lat}}$. BKZ reduction and sieving are then applied to the lattice basis built from $\mathbf{A}_{\text{lat}}$ as in Eq. (1), and the list $\mathcal{W}$ of short vectors is output.

**`Modulus_Switch_with_List_Decode`.** Given the list of short vectors $\mathcal{W} = \{\mathbf{w}_i = (\mathbf{x}_i^\top, \mathbf{y}_i^\top)^\top\}_{i=1}^N$ obtained from the sample step, the rounded vectors $\lfloor \frac{p}{q}\mathbf{A}_{\text{fft}}^\top \mathbf{x}_i \rceil$ are first computed for each $i$. A polar code $C(\widetilde{\mathbf{G}})$ over $\mathbb{Z}_p$ is then constructed, and the list decoding procedure from Lemma 2 is applied to obtain the corresponding information vectors $\{\widetilde{\mathbf{u}_i}\}$. Since $\widetilde{\mathbf{s}_{\text{enu}}} = \mathbf{0}$ is assumed, the term $\mathbf{A}_{\text{enu}}\widetilde{\mathbf{s}_{\text{enu}}}$ is omitted. Finally, this subroutine returns the list $\mathscr{L} = \{(\widetilde{\mathbf{u}_i}^\top, \ \frac{p}{q}\langle \mathbf{x}_i, \mathbf{b}\rangle)^\top\}_{i=1}^N$.

**`Solve_LWE_with_FFT`.** This subroutine computes, via FFT, the scores for all candidates $\widetilde{\mathbf{G}}^\top \widetilde{\mathbf{s}_{\text{fft}}}$ simultaneously. For a fixed guess of $\widetilde{\mathbf{s}_{\text{enu}}}$, the output $V$ is the maximum score

over all guesses of $\widetilde{\mathbf{G}}^\top \widetilde{\mathbf{s}_{\text{fft}}}$. Specifically, for $\mathbf{z} \in \mathbb{Z}_p^{k_{\text{fft}}}$, the function

$$f_0(\widetilde{\mathbf{u}_i}) = \sum_{(\widetilde{\mathbf{u}_i}, b_i) \in \mathscr{L}} e^{\frac{2i\pi}{p} b_i}$$

is defined. The FFT of $f_0$ is then computed and its real part is taken. It is readily verified that $\widehat{f_0}(\mathbf{z}) = F_{\mathcal{W}}^{(\text{new})}(\mathbf{0}, \mathbf{z})$. Next, the threshold $T$ is set to be the expected score under the correct guess. The value $V$ is then compared with $T$; if $V \geq T$, the guess is accepted.

**`Sub_LWE_Solver.`** For the reduced LWE instance $([\mathbf{A}_{\text{fft}} \mid \mathbf{A}_{\text{lat}}], \mathbf{b})$, an LWE solver is used; its complexity is negligible compared to the reduction, and it recovers both $\mathbf{s}_{\text{lat}}$ and $\mathbf{s}_{\text{fft}}$.

### 3.2 Score Distribution under the Correct Guess

In this subsection, we analyze the score distribution of our new variant under the correct guess, following the unified framework of [10]. For the three score functions in Eqs. (2)-(4) described earlier, the model in [10] provides estimates of the expectation and variance of the score under the correct guess, under reasonable assumptions. To similarly predict the expectation of the score for our new variant $F_{\mathcal{W}}^{(\text{new})}$, we adopt analogous assumptions.

**Assumption 1.** *For $\mathbf{w}_i = (\mathbf{x}_i^\top, \mathbf{y}_i^\top)^\top \in \mathcal{W}$, the components of $\left\{\frac{p}{q}\mathbf{A}_{\text{fft}}^\top \mathbf{x}_i\right\}$ are mutually independent and each follows the uniform distribution over $\left[-\frac{1}{2}, \frac{1}{2}\right)$. Moreover, the three inner products*

$$\langle \mathbf{w}_i, \mathbf{e}' \rangle, \quad \left\langle \left\{\frac{p}{q}\mathbf{A}_{\text{fft}}^\top \mathbf{x}_i\right\}, \mathbf{s}_{\text{fft}} \right\rangle, \quad \langle \widetilde{\mathbf{e}_i}, \mathbf{s}_{\text{fft}} \rangle$$

*are assumed to be mutually independent.*

*Remark 1.* Let

$$\mathbf{r}_i \triangleq \left\{\frac{p}{q}\mathbf{A}_{\text{fft}}^\top \mathbf{x}_i\right\}, \qquad \widetilde{\mathbf{e}}_i \triangleq \left\lfloor \frac{p}{q}\mathbf{A}_{\text{fft}}^\top \mathbf{x}_i \right\rceil - \widetilde{\mathbf{G}}\widetilde{\mathbf{u}}_i.$$

Although both $\mathbf{r}_i$ and $\widetilde{\mathbf{e}}_i$ originate from the same vector $\frac{p}{q}\mathbf{A}_{\text{fft}}^\top \mathbf{x}_i$, they arise from different sources: $\mathbf{r}_i$ is the rounding error introduced by modulus switching, whereas $\widetilde{\mathbf{e}}_i$ is the decoding error introduced by lossy source coding. Therefore, within the modeling framework of Assumption 1, it is reasonable to treat them as independent error sources. Moreover, although both are multiplied by the same $\mathbf{s}_{\text{fft}}$ in the score function, their products can be treated as independent contributions in the subsequent expectation factorization. To see this at the second-moment level, consider each component and write

$$A = r_{i,j}, \qquad B = \widetilde{e}_{i,j}, \qquad C = (\mathbf{s}_{\text{fft}})_j.$$

Under the independence modeled in Assumption 1, and since $A$ and $B$ are zero-mean, we have

$$\begin{aligned} \text{Cov}(AC, BC) &= \mathbb{E}[ABC^2] - \mathbb{E}[AC]\mathbb{E}[BC] \\ &= \mathbb{E}[A]\mathbb{E}[B]\left(\mathbb{E}[C^2] - \mathbb{E}[C]^2\right) \\ &= \mathbb{E}[A]\mathbb{E}[B]\,\text{Var}(C) = 0. \end{aligned}$$

Thus the two products are uncorrelated in the second-moment sense, which provides additional justification for treating them as independent error contributions.

Strictly speaking, the covariance calculation only shows uncorrelatedness; the independence used later is part of Assumption 1, and this computation serves as a consistency check.

We now adapt an important lemma from [10, Lemma 6] to our setting, presenting a slightly more general version that will be used in our analysis.

**Lemma 3.** *(Adapted from [10, Lemma 6]) Let* $\mathbf{y} \in \mathbb{Z}_q^n$ *be a random vector whose components are independently drawn from a symmetric distribution with mean* 0 *and variance* $\sigma^2$. *Define* $\rho_q(\sigma, t) = e^{-\frac{2\pi^2\sigma^2t^2}{q^2}}$. *Then, for any fixed* $\mathbf{x} \in \mathbb{Z}_q^n$,

$$\mathbb{E}\left[\cos\left(\frac{2\pi}{q}\langle \mathbf{x}, \mathbf{y}\rangle\right)\right] \approx \rho_q(\sigma, \|\mathbf{x}\|),$$

*and*

$$\mathbb{V}\left[\cos\left(\frac{2\pi}{q}\langle \mathbf{x}, \mathbf{y}\rangle\right)\right] \approx \frac{1}{2} - \rho_q(\sigma, \|\mathbf{x}\|)^2 + \frac{1}{2}\rho_q(\sigma, \|\mathbf{x}\|)^4.$$

Following the approach of [10], we now estimate the expected score for our new variant. Importantly, this estimation only requires the expectation and does not rely on the problematic independence assumptions used in earlier work. Under the correct guess, the score function in Eq. (6) takes the form

$$F_{\mathcal{W}}^{(\mathrm{new})}(\mathbf{s}_{\mathrm{enu}}, \mathbf{s}_{\mathrm{fft}}) = \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left(\frac{2\pi}{q}\langle \mathbf{x}_i, \mathbf{e}'\rangle + \frac{2\pi}{p}\left\langle \left\{\frac{p}{q}\mathbf{A}_{\mathrm{fft}}^\top \mathbf{x}_i\right\}, \mathbf{s}_{\mathrm{fft}}\right\rangle + \frac{2\pi}{p}\langle \widetilde{\mathbf{e}}_i, \mathbf{s}_{\mathrm{fft}}\rangle\right).$$

Taking expectations on both sides and applying Lemma 1 together with Assumption 1, the expectation of the score function factors into a product of three independent expectations:

$$\begin{aligned}\mathbb{E}\left[F_{\mathcal{W}}^{(\mathrm{new})}(\mathbf{s}_{\mathrm{enu}}, \mathbf{s}_{\mathrm{fft}})\right] = \sum_{\mathbf{w}_i \in \mathcal{W}} &\mathbb{E}\left[\cos\left(\frac{2\pi}{q}\langle \mathbf{x}_i, \mathbf{e}'\rangle\right)\right] \cdot \mathbb{E}\left[\cos\left(\frac{2\pi}{p}\langle \widetilde{\mathbf{e}}_i, \mathbf{s}_{\mathrm{fft}}\rangle\right)\right] \\ &\cdot \mathbb{E}\left[\cos\left(\frac{2\pi}{p}\left\langle \left\{\frac{p}{q}\mathbf{A}_{\mathrm{fft}}^\top \mathbf{x}_i\right\}, \mathbf{s}_{\mathrm{fft}}\right\rangle\right)\right].\end{aligned}$$

For the second term, we approximate the norm of $\widetilde{\mathbf{e}}_i$ by the Gilbert-Varshamov type bound in Eq. (5), i.e. $\|\widetilde{\mathbf{e}}_i\| \approx p^{1-\frac{k_{\mathrm{fft}}}{n_{\mathrm{fft}}}}\sqrt{\frac{n_{\mathrm{fft}}}{2\pi e}}$. Then, applying Lemma 3 with $\mathbf{x} = \widetilde{\mathbf{e}}_i$ and $\mathbf{y} = \mathbf{s}_{\mathrm{fft}}$, we obtain

$$\mathbb{E}\left[\cos\left(\frac{2\pi}{p}\langle \widetilde{\mathbf{e}}_i, \mathbf{s}_{\mathrm{fft}}\rangle\right)\right] \approx \rho_p\left(\sigma_e,\ p^{1-\frac{k_{\mathrm{fft}}}{n_{\mathrm{fft}}}}\sqrt{\frac{n_{\mathrm{fft}}}{2\pi e}}\right).$$

By Lemma 1, the third expectation satisfies

$$\mathbb{E}\left[\cos\left(\frac{2\pi}{p}\left\langle \left\{\frac{p}{q}\mathbf{A}_{\mathrm{fft}}^\top \mathbf{x}_i\right\}, \mathbf{s}_{\mathrm{fft}}\right\rangle\right)\right] \approx \left(\sum_{t\in\mathbb{Z}} \Pr(\chi_e = t) \cdot \frac{p\sin(\pi t/p)}{\pi t}\right)^{n_{\mathrm{fft}}}.$$

Indeed, the components of $\mathbf{s}_{\text{fft}}$ and $\left\{\frac{p}{q}\mathbf{A}_{\text{fft}}^{\top}\mathbf{x}_i\right\}$ are mutually independent (the latter by Assumption 1). Hence, the products $\left\{\frac{p}{q}\mathbf{A}_{\text{fft}}^{\top}\mathbf{x}_i\right\}_j \cdot (\mathbf{s}_{\text{fft}})_j$ are mutually independent as well, where the subscript $j$ denotes the $j$-th component of the corresponding vector. For each component, assuming $\chi_e$ is discrete (otherwise replace the sum by an integral), conditional probability and Lemma 1 give

$$\mathbb{E}\left[\cos\left(\frac{2\pi}{p}\cdot\left\{\frac{p}{q}\mathbf{A}_{\text{fft}}^{\top}\mathbf{x}_i\right\}_j\cdot(\mathbf{s}_{\text{fft}})_j\right)\right] = \sum_{t\in\mathbb{Z}}\Pr(\chi_e=t)\cdot\frac{2\sin(\pi t/p)}{2\pi t/p}.$$

Applying Lemma 1 again to combine the $n_{\text{fft}}$ independent components yields the claimed expression.

Combining the three estimates above, we obtain the following approximation for the expected score under the correct guess:

$$\mathbb{E}\left[F_{\mathcal{W}}^{(\text{new})}\right] \approx \sum_{\mathbf{w}_i\in\mathcal{W}}\rho_q(\sigma_e,\|\mathbf{w}_i\|)\cdot\rho_p\left(\sigma_e,\, p^{1-\frac{k_{\text{fft}}}{n_{\text{fft}}}}\sqrt{\frac{n_{\text{fft}}}{2\pi e}}\right)\cdot\left(\sum_{t\in\mathbb{Z}}\Pr(\chi_e=t)\frac{p\sin(\pi t/p)}{\pi t}\right)^{n_{\text{fft}}},$$

where $\rho_q(\sigma,t) = e^{\frac{-2\pi^2\sigma^2t^2}{q^2}}$ is as defined in Lemma 3.

### 3.3 Experimental Validation on Small-Scale Instances

To validate the accuracy of our estimates, we conducted experiments on small-scale instances.[2] Following the KYBER-512 parameter settings, we set $q = 3329$ and $\chi_e = \mathcal{B}_3$. The other parameters were chosen as $n_{\text{enu}} = 5$, $n_{\text{fft}} = 10$, $k_{\text{fft}} = 3$, and $n_{\text{lat}} \in \{90, 100, 110, 120\}$, with $m = n = n_{\text{lat}} + n_{\text{enu}} + n_{\text{fft}}$. To obtain the short vector list $\mathcal{W}$, we first applied the progressive BKZ algorithm from the BLASter library [8] to reduce the lattice basis $\mathbf{B}$ in Eq. (1), and then used the progressive sieving routine from the G6K library [1] to output short vectors. The parameters for lattice reduction and sieving were fixed as $\beta_{\text{bkz}} = 60$, $\beta_{\text{sieve}} = 85$, $r_{\text{sat}} = \sqrt{4/3}$, and $f_{\text{sat}} = 0.9$, yielding $|\mathcal{W}| = \frac{1}{2}\cdot f_{\text{sat}}\cdot r_{\text{sat}}^{\beta_{\text{sieve}}} = 91857$.

We ran experiments for both the LSC-MS and MS-LSC routes. For rapid testing, we set $p = 256$. The results are summarised in Table 1. For the LSC-MS route, the expected score under the correct guess is indeed very small, which confirms our earlier conclusion that this route is not viable. In contrast, for the MS-LSC route, the observed expected scores are in excellent agreement with our theoretical predictions. This clearly demonstrates the soundness of our analysis and establishes the correctness of the proposed variant.

## 4 Complexity Estimation

In this section, we evaluate the complexity of our new variant following the estimation framework of [5]. The overall time complexity consists of three parts: (i) the cost of

[2] The source code for the small-scale experiments, FFT operation counts, and complexity estimates is available at https://github.com/yechenli-ux/MS-LSC.

**Table 1.** Small-scale experimental results for the MS-LSC and LSC-MS routes.

| Route | $n_{\text{lat}}$ | Actual expectation | Predicted expectation | Ratio |
|---|---|---|---|---|
| MS-LSC | 90 | 25765.72 | 25615.24 | 1.005 |
| | 100 | 17683.92 | 17596.85 | 1.004 |
| | 110 | 9793.03 | 9618.00 | 1.018 |
| | 120 | 3738.53 | 3732.07 | 1.002 |
| LSC-MS | 90 | 3.489 | – | – |
| | 100 | 1.242 | – | – |

lattice reduction and sieving, denoted $T_{\text{sample}}$, given by [5, Lemma 4.1]; (ii) the decoding cost $|\mathcal{W}| \cdot T_{\text{dec}}$, where $T_{\text{dec}}$ is given by Lemma 2; and (iii) the FFT cost $T_{\text{FFT}}$, given by Proposition 1. Both $T_{\text{dec}}$ and $T_{\text{FFT}}$ depend on the quantities $N_{\text{FFT}}^{(\text{add})}$ and $N_{\text{FFT}}^{(\text{mul})}$. Following the method of [5], we compute these values for $p = 512, 1024, 2048, 3072$. To highlight the efficiency gain when $p$ is a power of two or a product of small primes, we also include the cases $p = 513, 1025, 2049, 3071, 3329, 3330$. The results are summarized in Table 2. As can be seen, when $p$ is a power of two or factors into small primes, the corresponding values of $N_{\text{FFT}}^{(\text{add})}$ and $N_{\text{FFT}}^{(\text{mul})}$ are significantly smaller, leading to improved efficiency in both decoding and FFT.

**Table 2.** Values of $N_{\text{FFT}}^{(\text{add})}$ and $N_{\text{FFT}}^{(\text{mul})}$ for different $p$.

| $p$ | $N_{\text{FFT}}^{(\text{add})}$ | $N_{\text{FFT}}^{(\text{mul})}$ | $p$ | $N_{\text{FFT}}^{(\text{add})}$ | $N_{\text{FFT}}^{(\text{mul})}$ |
|---|---|---|---|---|---|
| 512 | 11648 | 4352 | 2049 | 272696 | 153876 |
| 513 | 19386 | 14904 | 3071 | 398274 | 379972 |
| 1024 | 26112 | 9728 | 3072 | 91264 | 36096 |
| 1025 | 61400 | 51480 | 3329 | 241540 | 118016 |
| 2048 | 57600 | 22016 | 3330 | 199248 | 157074 |

The total complexity of our variant for the parameters of KYBER is then evaluated using

$$T_{\text{total}} = T_{\text{sample}} + R\big(|\mathcal{W}| \cdot T_{\text{dec}} + T_{\text{FFT}}\big),$$

where $R$ denotes the number of iterations performed in the attack. Its value is determined by the target success probability; due to space limitations, we refer the reader to [5, Lemma 3.2] for the detailed computation.

In the parameter search, the score distribution for incorrect guesses is modelled as a normal distribution $\mathcal{N}(0, |\mathcal{W}|/2)$. Although both Ducas [7] and Carrier [5] point out that the tail probability of the score distribution for incorrect guesses can be larger than that predicted by the normal model, it is observed that under the optimal parameters for KYBER given in [5], the normal prediction is always conservative (i.e. overestimates the tail probability). This justifies the use of the normal distribution as a safe approximation. For consistency, the optimal complexities from [5] are not directly adopted; instead, the complexity of the dual attack with lossy source coding alone is recomputed under the framework of [10], which yields results very close to those reported in [5]. Tables 3 and 4 report the total complexities of the new variant for $p = 512, 1024, 2048, 3072$, as

**Table 3.** Total complexity (in bits) of our variant for KYBER under different cost models and intermediate moduli $p$.

| Scheme | Model | $p = 512$ | $p = 1024$ | $p = 2048$ | $p = 3072$ | model in [5] |
|---|---|---|---|---|---|---|
| KYBER-512 | C0 | **118.09** | 118.76 | 120.02 | 120.46 | 121.82 |
| | CC | 139.27 | **139.20** | 139.33 | 139.51 | 139.40 |
| | CN | 134.70 | 134.75 | 134.61 | **134.42** | 134.54 |
| KYBER-768 | C0 | **172.36** | 172.83 | 172.48 | 172.39 | 172.75 |
| | CC | 194.75 | 194.60 | 195.03 | **194.02** | 195.09 |
| | CN | **189.36** | 189.62 | 189.53 | 189.87 | 189.63 |
| KYBER-1024 | C0 | 238.40 | **238.12** | 238.40 | 238.40 | 238.42 |
| | CC | 259.59 | 259.65 | 259.65 | **259.40** | 259.44 |
| | CN | **253.68** | 253.78 | 254.00 | 253.80 | 254.19 |

well as for the lossy-source-coding-only case, along with the corresponding values of $R \cdot T_{\mathrm{FFT}}$ and $R \cdot |\mathcal{W}| T_{\mathrm{dec}}$.

The models C0, CC, and CN refer to different cost models for lattice reduction, following exactly the same frameworks as in [5]. All three use the sieving oracle of [4] and are consistent with those presented in [2].

- **C0**: Core-SVP model [3]; assumes a single SVP call suffices and neglects lower-order exponent terms.
- **CC**: Classical circuit model [12]; the most detailed estimate to date, but ignores memory access costs and may underestimate true complexity.
- **CN**: Classical query model [4]; based on a naive nearest-neighbour search.

From the data, we observe that our new variant can achieve a better total complexity than [5] in all cases (marked in bold in the tables), although the improvements are often modest. At the same time, our variant offers greater flexibility in parameter selection. More importantly, in terms of the FFT cost $R \cdot T_{\mathrm{FFT}}$ and the decoding cost $R \cdot |\mathcal{W}| T_{\mathrm{dec}}$, our variant achieves improvements of 1–6 bits and 2–7 bits, respectively, in most parameter settings. The only exception is the CN model for KYBER-1024, where the lossy-source-coding-only model achieves the optimal FFT complexity due to the relatively small choice of $k_{\mathrm{fft}}$; the detailed parameter choices are provided in Appendix B.

## 5 Conclusion

In this paper, we propose a new dual-sieve-FFT variant integrating modulus switching and lossy source coding. We show that only the MS-LSC route is viable, reducing the FFT enumeration from $q^{n_{\mathrm{fft}}}$ to $p^{k_{\mathrm{fft}}}$ and offering greater flexibility in parameter selection. Using the framework of [10], we analyze the score expectation under the correct guess and validate our estimates via small-scale experiments. When applied to KYBER, our variant achieves a total complexity comparable to or better than that of the work of Carrier et al. [5], while reducing the FFT and decoding costs by 1–6 bits and 2–7 bits, respectively. These improvements, though modest, are meaningful in practice, especially

**Table 4.** FFT and decoding complexities multiplied by $R$ (in bits) for KYBER under different cost models and intermediate moduli $p$.

| Scheme | Model | $R \cdot T_{\mathrm{FFT}}$ | | | | |
|---|---|---|---|---|---|---|
| | | $p = 512$ | $p = 1024$ | $p = 2048$ | $p = 3072$ | model in [5] |
| KYBER-512 | C0 | **109.75** | 109.76 | 107.76 | 112.55 | 114.97 |
| | CC | 134.57 | 136.60 | **125.48** | 130.86 | 133.39 |
| | CN | 131.23 | 131.61 | **125.48** | 130.86 | 131.71 |
| KYBER-768 | C0 | 169.67 | 168.47 | 170.11 | **167.55** | 169.01 |
| | CC | 192.08 | **190.07** | 191.52 | 192.35 | 192.64 |
| | CN | 184.42 | 187.24 | 184.05 | **180.66** | 185.08 |
| KYBER-1024 | C0 | 229.79 | 231.82 | 229.83 | **228.81** | 232.27 |
| | CC | 256.98 | **253.36** | 253.39 | 254.97 | 257.24 |
| | CN | 249.33 | 250.53 | 251.98 | 250.72 | **245.46** |
| Scheme | Model | $R \cdot \lvert\mathcal{W}\rvert T_{\mathrm{dec}}$ | | | | |
| | | $p = 512$ | $p = 1024$ | $p = 2048$ | $p = 3072$ | model in [5] |
| KYBER-512 | C0 | **117.57** | 118.53 | 119.90 | 120.39 | 121.79 |
| | CC | **122.25** | 124.04 | 125.21 | 125.70 | 127.31 |
| | CN | 119.54 | **118.22** | 125.21 | 125.70 | 125.63 |
| KYBER-768 | C0 | **163.67** | 168.08 | 167.41 | 169.53 | 169.92 |
| | CC | 165.21 | 166.37 | **164.92** | 168.45 | 167.23 |
| | CN | 167.05 | **163.96** | 168.16 | 169.07 | 171.89 |
| KYBER-1024 | C0 | **216.44** | 217.39 | 218.77 | 219.46 | 221.07 |
| | CC | **213.27** | 216.05 | 217.22 | 219.13 | 219.11 |
| | CN | 213.90 | **213.43** | 216.02 | 216.29 | 218.96 |

in multi-target settings where the FFT and decoding steps are amortized over multiple targets.

Our results further confirm that the sampling step remains the main bottleneck in the total complexity, which is consistent with the findings of Carrier et al. To significantly reduce the overall cost, one would need a technique that lowers the FFT complexity with smaller introduced error — this is precisely why the lossy-source-coding approach of Carrier et al. outperforms MATZOV. Although the results of Carrier et al. are already quite strong, our variant retains its own merit by enabling a wider range of parameter trade-offs, which may be beneficial in broader application scenarios. Our results also deepen the understanding of the dual-sieve-FFT attack. In particular, although the parameters of the three main steps are interdependent, our findings indicate that meaningful improvements can still be achieved by refining individual steps. A promising direction for future work is to reduce the sampling cost by replacing the high-dimensional sieving with multiple rounds of lower-dimensional lattice reduction and sieving, while ensuring that the required number of short vectors is maintained. This approach, similar in spirit to MATZOV, may offer a potential avenue for lowering the sampling complexity and achieving a more balanced overall attack.

**Acknowledgments.** We would like to thank the anonymous reviewers for their constructive comments. This study was supported by the National Key Research and Development Program of China (No. 2024YFB4504700) and the National Natural Science Foundation of China (No. 12371526).

## A LSC-MS Route

We now examine the LSC-MS route, where lossy source coding is applied before modulus switching. When lossy source coding is applied first, the score function becomes

$$\begin{aligned} F^1_{\mathcal{W}}(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left( \frac{2\pi}{q} \langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}} \widetilde{\mathbf{s}_{\text{enu}}} \rangle - \frac{2\pi}{q} \langle \mathbf{G}\mathbf{u}_i, \widetilde{\mathbf{s}_{\text{fft}}} \rangle \right) \\ &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left( \frac{2\pi}{q} \langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}} \widetilde{\mathbf{s}_{\text{enu}}} \rangle - \frac{2\pi}{q} \langle \mathbf{u}_i, \mathbf{G}^\top \widetilde{\mathbf{s}_{\text{fft}}} \rangle \right). \end{aligned}$$

Then modulus switching is applied, yielding

$$F^2_{\mathcal{W}}(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) = \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left( \frac{2\pi}{q} \langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}} \widetilde{\mathbf{s}_{\text{enu}}} \rangle - \frac{2\pi}{p} \langle \left\lfloor \frac{p}{q} \mathbf{u}_i \right\rceil, \mathbf{G}^\top \widetilde{\mathbf{s}_{\text{fft}}} \rangle \right).$$

Under the correct guess $(\widetilde{\mathbf{s}_{\text{enu}}}, \widetilde{\mathbf{s}_{\text{fft}}}) = (\mathbf{s}_{\text{enu}}, \mathbf{s}_{\text{fft}})$, we have

$$\begin{aligned} F^2_{\mathcal{W}} &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left( \frac{2\pi}{q} \left( \langle \mathbf{x}_i, \mathbf{b} - \mathbf{A}_{\text{enu}} \mathbf{s}_{\text{enu}} \rangle - \langle \mathbf{u}_i, \mathbf{G}^\top \mathbf{s}_{\text{fft}} \rangle \right) + \frac{2\pi}{p} \left\langle \left\{ \frac{p}{q} \mathbf{u}_i \right\}, \mathbf{G}^\top \mathbf{s}_{\text{fft}} \right\rangle \right) \\ &= \sum_{\mathbf{w}_i \in \mathcal{W}} \cos\left( \frac{2\pi}{q} \left( \langle \mathbf{w}_i, \mathbf{e}' \rangle + \langle \mathbf{e}_i, \mathbf{s}_{\text{fft}} \rangle \right) + \frac{2\pi}{p} \left\langle \left\{ \frac{p}{q} \mathbf{u}_i \right\}, \mathbf{G}^\top \mathbf{s}_{\text{fft}} \right\rangle \right), \end{aligned}$$

where $\mathbf{e}_i = \mathbf{A}_{\text{fft}}^\top \mathbf{x}_i - \mathbf{G}\mathbf{u}_i$. Here, $\mathbf{G}^\top \mathbf{s}_{\text{fft}}$ is not necessarily a short vector; consequently, the term $\left\langle \left\{ \frac{p}{q} \mathbf{u}_i \right\}, \mathbf{G}^\top \mathbf{s}_{\text{fft}} \right\rangle$ cannot be guaranteed to be small. This introduces an uncontrollable error term that significantly reduces the expectation of $F^2_{\mathcal{W}}$ under the correct guess, thereby compromising the distinguishability between correct and incorrect guesses. Moreover, unlike the MS-LSC route, this order does not offer any reduction in decoding complexity: since $\mathbf{G}$ remains defined over $\mathbb{Z}_q$, the decoding step must still be performed over the original modulus, and the efficiency gain from working in a smaller field is lost. Consequently, the LSC-MS route fails to provide the desired performance improvement and is therefore not feasible for our purposes.

## B Parameter Details for Optimal Complexities

We provide in the following tables the detailed parameters corresponding to the optimal complexities. For each scheme, cost model, and intermediate modulus $p$, we list the parameters $(m, n_{\text{enu}}, n_{\text{fft}}, k_{\text{fft}}, \beta_{\text{bkz}}, \beta_{\text{sieve}}, R, T_{\text{sample}})$ that achieve the best total complexity.

**Table 5.** Optimal parameters for $p = 512$ under different schemes and models.

| Scheme | Model | $m$ | $n_{\text{enu}}$ | $n_{\text{fft}}$ | $k_{\text{fft}}$ | $\beta_{\text{bkz}}$ | $\beta_{\text{sieve}}$ | $\log_2 R$ | $T_{\text{sample}}$ |
|---|---|---|---|---|---|---|---|---|---|
| KYBER-512 | C0 | 468 | 1 | 37 | 10 | 395 | 395 | 2.84 | 116.34 |
| | CC | 459 | 5 | 49 | 12 | 383 | 386 | 9.39 | 139.21 |
| | CN | 447 | 3 | 49 | 12 | 385 | 389 | 6.05 | 134.57 |
| KYBER-768 | C0 | 628 | 5 | 80 | 16 | 586 | 586 | 8.08 | 172.11 |
| | CC | 616 | 8 | 91 | 18 | 580 | 573 | 12.32 | 194.51 |
| | CN | 651 | 9 | 90 | 17 | 582 | 575 | 13.74 | 189.31 |
| KYBER-1024 | C0 | 820 | 9 | 102 | 22 | 813 | 813 | 13.74 | 238.40 |
| | CC | 787 | 9 | 131 | 25 | 811 | 792 | 13.74 | 259.34 |
| | CN | 810 | 10 | 126 | 24 | 813 | 794 | 15.15 | 253.61 |

**Table 6.** Optimal parameters for $p = 1024$ under different schemes and models.

| Scheme | Model | $m$ | $n_{\text{enu}}$ | $n_{\text{fft}}$ | $k_{\text{fft}}$ | $\beta_{\text{bkz}}$ | $\beta_{\text{sieve}}$ | $\log_2 R$ | $T_{\text{sample}}$ |
|---|---|---|---|---|---|---|---|---|---|
| KYBER-512 | C0 | 488 | 1 | 39 | 9 | 394 | 394 | 2.84 | 116.05 |
| | CC | 470 | 5 | 50 | 11 | 382 | 385 | 9.39 | 138.94 |
| | CN | 454 | 2 | 50 | 11 | 385 | 389 | 4.40 | 134.57 |
| KYBER-768 | C0 | 666 | 7 | 65 | 14 | 588 | 588 | 10.91 | 172.70 |
| | CC | 647 | 8 | 89 | 16 | 580 | 573 | 12.32 | 194.54 |
| | CN | 643 | 6 | 90 | 16 | 582 | 575 | 9.49 | 189.31 |
| KYBER-1024 | C0 | 818 | 9 | 103 | 20 | 812 | 812 | 13.74 | 238.10 |
| | CC | 790 | 10 | 129 | 22 | 812 | 793 | 15.15 | 259.63 |
| | CN | 828 | 8 | 128 | 22 | 813 | 794 | 12.32 | 253.62 |

**Table 7.** Optimal parameters for $p = 2048$ under different schemes and models.

| Scheme | Model | $m$ | $n_{\text{enu}}$ | $n_{\text{fft}}$ | $k_{\text{fft}}$ | $\beta_{\text{bkz}}$ | $\beta_{\text{sieve}}$ | $\log_2 R$ | $T_{\text{sample}}$ |
|---|---|---|---|---|---|---|---|---|---|
| KYBER-512 | C0 | 508 | 1 | 36 | 8 | 395 | 395 | 2.84 | 116.34 |
| | CC | 478 | 5 | 54 | 10 | 382 | 386 | 9.39 | 139.09 |
| | CN | 477 | 5 | 45 | 9 | 385 | 389 | 9.39 | 134.61 |
| KYBER-768 | C0 | 646 | 6 | 72 | 13 | 586 | 586 | 9.49 | 172.11 |
| | CC | 612 | 6 | 94 | 15 | 581 | 574 | 9.49 | 194.79 |
| | CN | 677 | 8 | 82 | 14 | 582 | 576 | 12.32 | 189.49 |
| KYBER-1024 | C0 | 820 | 9 | 102 | 18 | 813 | 813 | 13.74 | 238.40 |
| | CC | 790 | 10 | 129 | 20 | 812 | 793 | 15.15 | 259.63 |
| | CN | 791 | 9 | 130 | 20 | 813 | 794 | 13.74 | 253.59 |

**Table 8.** Optimal parameters for $p = 3072$ under different schemes and models.

| Scheme | Model | $m$ | $n_{\text{enu}}$ | $n_{\text{fft}}$ | $k_{\text{fft}}$ | $\beta_{\text{bkz}}$ | $\beta_{\text{sieve}}$ | $\log_2 R$ | $T_{\text{sample}}$ |
|---|---|---|---|---|---|---|---|---|---|
| KYBER-512 | C0 | 468 | 1 | 38 | 8 | 394 | 394 | 2.84 | 116.05 |
| | CC | 468 | 5 | 48 | 9 | 384 | 387 | 9.39 | 139.51 |
| | CN | 454 | 5 | 48 | 9 | 384 | 388 | 9.39 | 134.29 |
| KYBER-768 | C0 | 666 | 7 | 72 | 12 | 586 | 586 | 10.91 | 172.11 |
| | CC | 594 | 8 | 92 | 14 | 581 | 574 | 12.32 | 194.77 |
| | CN | 632 | 8 | 86 | 13 | 584 | 577 | 12.32 | 189.87 |
| KYBER-1024 | C0 | 819 | 9 | 103 | 17 | 813 | 813 | 13.74 | 238.40 |
| | CC | 808 | 11 | 135 | 19 | 811 | 792 | 16.57 | 259.35 |
| | CN | 826 | 8 | 129 | 19 | 813 | 794 | 12.32 | 253.62 |

**Table 9.** Optimal parameters for the lossy-source-coding-only attack under different schemes and models.

| Scheme | Model | $m$ | $n_{\text{enu}}$ | $n_{\text{fft}}$ | $k_{\text{fft}}$ | $\beta_{\text{bkz}}$ | $\beta_{\text{sieve}}$ | $\log_2 R$ | $T_{\text{sample}}$ |
|---|---|---|---|---|---|---|---|---|---|
| KYBER-512 | C0 | 488 | 1 | 39 | 8 | 393 | 393 | 2.84 | 115.76 |
| | CC | 472 | 5 | 47 | 9 | 383 | 387 | 9.39 | 139.38 |
| | CN | 470 | 4 | 47 | 9 | 384 | 388 | 7.71 | 134.31 |
| KYBER-768 | C0 | 646 | 6 | 71 | 12 | 587 | 587 | 9.49 | 172.40 |
| | CC | 624 | 6 | 93 | 14 | 581 | 574 | 9.49 | 194.80 |
| | CN | 624 | 9 | 90 | 13 | 583 | 576 | 13.74 | 189.57 |
| KYBER-1024 | C0 | 840 | 9 | 100 | 17 | 813 | 813 | 13.74 | 238.40 |
| | CC | 835 | 10 | 129 | 19 | 810 | 791 | 15.15 | 259.09 |
| | CN | 820 | 10 | 123 | 18 | 815 | 796 | 15.15 | 254.19 |

## References


1. Albrecht, M.R., Ducas, L., Herold, G., Kirshanova, E., Postlethwaite, E.W., Stevens, M.: The general sieve kernel and new records in lattice reduction. In: Ishai, Y., Rijmen, V. (eds.) Advances in Cryptology – EUROCRYPT 2019. pp. 717–746. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-030-17656-3_25
2. Albrecht, M.R., Shen, Y.: Quantum augmented dual attack. Cryptology ePrint Archive, Paper 2022/656 (2022), https://eprint.iacr.org/2022/656
3. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange: a new hope. In: 25th USENIX Security Symposium (USENIX Security 16). pp. 327–343. USENIX Association, Austin, TX (Aug 2016), https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim
4. 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. p. 10–24. SODA '16, Society for Industrial and Applied Mathematics, USA (2016), https://dl.acm.org/doi/10.5555/2884435.2884437
5. 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.) Advances in Cryptology – CRYPTO 2025. pp. 444–476. Springer Nature Switzerland, Cham (2025), https://doi.org/10.1007/978-3-032-01855-7_15
6. Ducas, L., Pulles, L.N.: Does the dual-sieve attack on learning with errors even work? In: Handschuh, H., Lysyanskaya, A. (eds.) Advances in Cryptology – CRYPTO 2023. pp. 37–69. Springer Nature Switzerland, Cham (2023), https://doi.org/10.1007/978-3-031-38548-3_2
7. Ducas, L., Pulles, L.N.: Accurate score prediction for dual-sieve attacks: L. ducas, ln pulles. Journal of Cryptology **39**(1), 8 (2026), https://doi.org/10.1007/s00145-025-09560-7
8. Ducas, L., Pulles, L.N., Stevens, M.: Towards a modern lll implementation. In: Hanaoka, G., Yang, B.Y. (eds.) Advances in Cryptology – ASIACRYPT 2025. pp. 65–99. Springer Nature Singapore, Singapore (2026). https://doi.org/10.1007/978-981-95-5099-9_3
9. Guo, Q., Johansson, T.: Faster dual lattice attacks for solving lwe with applications to crystals. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2021. pp. 33–62. Springer International Publishing, Cham (2021), https://doi.org/10.1007/978-3-030-92068-5_2
10. Li, Y., Zheng, Q.: Unified dual attack analyses: Covariance-based score distribution prediction for LWE. Cryptology ePrint Archive, Paper 2026/1048 (2026), https://eprint.iacr.org/2026/1048
11. 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
12. MATZOV, Q.: Report on the security of lwe: Improved dual lattice attack (2022)
13. Micciancio, D., Regev, O.: Lattice-based cryptography. In: Bernstein, D.J., Buchmann, J., Dahmen, E. (eds.) Post-Quantum Cryptography. pp. 147–191. Springer Berlin Heidelberg, Berlin, Heidelberg (2009). https://doi.org/10.1007/978-3-540-88702-7_5, https://doi.org/10.1007/978-3-540-88702-7_5
14. 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