Provable FFT-Accelerated Dual Attack on LWE Using the General Dual Lattice

Jiale Li[1], Li-Ping Wang[1], and Huaxiong Wang[2]

[1] Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China
{lijiale,wangliping}@iie.ac.cn
[2] School of Physical & Mathematical Sciences, Nanyang Technological University, Singapore
hxwang@ntu.edu.sg

Abstract. The dual attack is a central tool for evaluating the security of the Learning with Errors (LWE) problem, which is the foundation of the lattice-based cryptosystems such as ML-KEM. Since the introduction of FFT acceleration by Guo and Johansson (ASIACRYPT 2021) and subsequent optimizations in the MATZOV report, FFT-based dual attacks have achieved concrete security estimates for ML-KEM. However, these results rely on one specific independence heuristic that is known to fail in certain parameter regimes (CRYPTO 2023).

Recent work by Pouly and Shen (EUROCRYPT 2024) initiated the study of provable dual attacks through explicit geometric analysis, but restricted the analysis to the orthogonal lattice, requiring approximately $m \approx 2n$ LWE samples. Qu and Xu (ASIACRYPT 2025) subsequently introduced modulus switching into the provable dual-attack setting, enabling evaluation over a smaller modulus, with a revised analysis given in their later version. More recently, a closely related ePrint revisited the concrete complexity estimates of these provable attacks after correcting the norm-bound formula used in the earlier analysis.

In this work, we provide a provable correctness analysis of the Guo–Johansson-style FFT-based dual attack over the general dual lattice. By extending the analysis from the orthogonal lattice to $L_q(\mathbf{A}^T)$, where $\mathbf{y}$ is not restricted to $\mathbf{0}$, we reduce the required number of LWE samples from $m \approx 2n$ to $m \approx n$ while retaining the computational efficiency of FFT-based score evaluation. We further derive an explicit upper bound on the false-positive probability via an integral estimate for a geometric distinguishing event.

For concrete complexity estimation, we follow the abstractions and cost models used in prior provable analyses while refining the probability-dependent treatment of the norm bounds. We obtain attack costs of 216, 317, and 426 bits for ML-KEM-512, ML-KEM-768, and ML-KEM-1024, respectively. Compared with the recomputed estimates of the previous best published provable dual attack, these represent reductions of 22, 30, and 52 bits, respectively.

Keywords: LWE · Provable dual attack · Lattices · FFT · False positive probability · Probability-decreasing enumeration.

1 Introduction

The Learning with Errors (LWE) problem, introduced by Regev in 2005 [25], has become a cornerstone of lattice cryptography. In the security analysis of ML-KEM, the dual attack — especially after the FFT acceleration introduced by Guo and Johansson [13] and the MATZOV report [18] — is widely recognized as one of the most effective attacks, with estimated costs surpassing those of the primal attack and even questioning their claimed security level. This has made the dual attack a central tool for evaluating lattice-based schemes and spurred intense research.

The original idea of the dual attack was proposed in [14] by using short dual vectors to separate points close to a lattice from those far away. Aharonov and Regev [1] formalized this method in 2004 and clarified its connection to the decisional Bounded Distance Decoding (BDD) problem. Later, the dual idea was brought into LWE analysis and formally established the basic dual attack paradigm [19] . Given an LWE instance $(\mathbf{A}, \mathbf{b} = \mathbf{As} + \mathbf{e} \mod q)$, an attacker finds short dual vectors $\mathbf{w}$ satisfying $\mathbf{w}^T\mathbf{A} \equiv 0 \mod q$. Then the inner product $\mathbf{w}^T\mathbf{b} = \langle \mathbf{w}, \mathbf{e} \rangle \mod q$ has an absolute value much smaller than $q/2$ since the error $\mathbf{e}$ is small. In contrast, for uniformly random $\mathbf{b}$, the inner product $\mathbf{w}^T\mathbf{b}$ is nearly uniform modulo q. This difference allows one to distinguish LWE instances from random ones.

However, this basic dual attack is hardly practical. The main obstacles are twofold. First, the attack needs a distinct short vector for each LWE sample, and computing short vectors via lattice reduction is very expensive. Second, it requires an exponential number of LWE samples, which is unrealistic. Alkim et al. [3] made a big advance by introducing lattice sieving into the dual attack. Lattice sieving not only finds the shortest vectors in the dual lattice, but more importantly, it generates exponentially many short vectors in a single run so that the information of these exponentially many short dual vectors can be leveraged to improve a distinguisher, for example, by summing a score function over them [16]. This greatly amortizes the cost of short vector sampling.

Another milestone is the use of FFT to compute the score function for many candidates simultaneously in [17]. Guo and Johansson [13] successfully applied this idea to the dual attack, significantly reducing the cost of the guessing phase. Furthermore, [18] optimized the FFT computation. These two works demonstrated the practical potential of the dual attack for moderate-dimensional LWE instances and thereby had a substantial impact on the concrete security strength of NIST candidates such as ML-KEM.

Despite these efficiency improvements, the correctness analysis of the above attacks relies heavily on a specific "independence heuristic". This heuristic assumes that short vectors obtained from lattice sieving or reduction are mutually independent, and that the scores of different secret guesses are independent. However, [10,11] found experimentally that theoretical predictions based on this heuristic deviate significantly from actual attack performance. This strongly suggests that there exist certain correlations among score terms, casting serious doubt on the validity of the independence heuristic.

The correlations arise from two distinct sources. First, the short dual vectors produced by sieving are inherently correlated, which implies statistical dependence among their score contributions. Second, all secret guesses are evaluated on the same LWE instance $(\mathbf{A}, \mathbf{b})$, introducing structural correlations among their scores. Later, [4] shows theoretically that the score terms have non-zero correlation when dual vectors are not mutually independent by explicitly computing the covariance of score terms.

To eliminate reliance on statistical assumptions, the first provable dual attack was proposed in [21]. Instead of hypothesis testing, the algorithm uses the geometric distance between lattice points. Its analysis is restricted to the orthogonal lattice $L_q^{\perp}(\mathbf{A}_{\text{dual}})$. This leads to a required sample size of about $m \approx 2n$, and the computational complexity still has a big gap compared to heuristic attacks.

Subsequently, modulus switching was introduced in [24], enabling FFT-based evaluation over a smaller modulus and further reducing the complexity of provable attacks. Yet their analysis remains confined to the orthogonal lattice, so the sample size requirement is not fundamentally improved. The revised analysis in [23] provides updated complexity estimates for the provable FFT-based dual attack. More recently, [27] recomputed the complexity estimates of [21] and [23] after correcting the norm-bound formula in [21].

Apart from improvements in the correctness analysis of the distinguishing algorithm, there have also been advances in guessing strategies for the dual attack. In [18], the probability-decreasing key enumeration strategy was employed, but this estimate was later pointed out to lack theoretical justification [10], and the gap was subsequently filled by a rigorous theoretical proof [12,6,5]. Furthermore, the probability-decreasing key enumeration strategy with early termination was proposed [5]. A different guessing paradigm was introduced in [7,8], where the enumeration step is replaced by assuming the secret in the guess window is zero and iterating this assumption over various choices for the enumeration part.

Quantum acceleration has also been introduced into provable dual attacks. Quantum mean estimation and quantum search have been applied to speed up secret enumeration and FFT threshold detection [2]. Recently, quantum rejection sampling has been employed to propose new provable quantum dual attack algorithms [29,9]. Concurrently, exploiting algebraic structures of Module-/Ring-LWE, it was shown that symmetries in power-of-two cyclotomic rings can improve the hit probability of hybrid guessing strategies without altering the distinguishing analysis [15,20]; experimental results indicate that such structure-aware hybrid dual attacks can reduce the estimated security of ML-KEM by several bits under standard cost models [20].

We provide a provable correctness analysis of the Guo-Johansson-style FFT-based dual attack over the general dual lattice. Our framework extends the provable analysis from the orthogonal lattice to the general dual lattice, reducing the required number of LWE samples from $m \approx 2n$ to $m \approx n$, and provides a rigorous analysis for the FFT-accelerated score function. We re-examine the geometric nature of the dual attack and, using rigorous integral estimation techniques, derive an explicit upper bound on the false positive probability.

1.1 Our Contributions

This paper incorporates the attack structure of [13] into a provable framework pioneered by [21] and further developed in [23]. Building on these works, we retain the general abstractions for short-vector generation and complexity estimation, while extending the provable analysis to the general dual lattice. In particular, we provide a rigorous analysis of the FFT score function and geometric separation conditions, and refine the probabilistic treatment of the norm bounds under the actual CBD distribution used in ML-KEM. The resulting framework combines the sample efficiency of the general dual lattice with the computational efficiency of FFT acceleration while providing provable correctness guarantees.

- **Provable correctness analysis of an FFT-based dual attack over the general dual lattice.** Prior provable dual attacks [21,22,23] are confined to the orthogonal lattice

$$L_q^{\perp}(\mathbf{A}_{\text{dual}}) = \left\{\mathbf{w} \in \mathbb{Z}^m : \mathbf{w}^T\mathbf{A}_{\text{dual}} = 0 \pmod q\right\},$$

  where $\mathbf{A}_{\text{dual}}$ is defined from $\mathbf{A}$ in Section 3.2. This restriction in prior provable frameworks requires approximately $m \approx 2n$ LWE samples. Building on the FFT-based attack structure of [13], we extend the correctness analysis to the general dual lattice

$$L_q(\mathbf{A}_{\text{dual}}^T) = \left\{(\mathbf{w}, \mathbf{y}) : \mathbf{w}^T\mathbf{A}_{\text{dual}} = \mathbf{y}^T \pmod q\right\},$$

  reducing the sample requirement from $m \approx 2n$ to $m \approx n$. We retain the FFT score function of [13] and provide a rigorous analysis of its behavior in the general dual-lattice setting. This combines the sample efficiency of the general dual lattice with the computational efficiency of FFT-based score evaluation within the same correctness-analysis framework.
- **Rigorous probabilistic analysis of attack correctness.** Even with a well-concentrated score function, the attack can still fail if an incorrect guess happens to yield a score close to that of the correct one. We derive an explicit upper bound on the false-positive probability via an integral estimate for the corresponding geometric distinguishing event. In addition, for the norm bound used in our analysis, we derive a rigorous high-probability bound directly under the centered binomial distribution used in ML-KEM, with its associated probability explicitly quantified. These results provide rigorous probabilistic guarantees for the correctness analysis of the proposed attack.
- **Concrete complexity evaluation.** We instantiate the proposed framework for the ML-KEM parameter sets and provide concrete attack-cost estimates under a common probabilistic criterion for the norm bounds. Under this criterion, our attack achieves costs of 216, 317, and 426 bits for ML-KEM-512, ML-KEM-768, and ML-KEM-1024, respectively, improving upon the best previously published provable dual attack of [21] by 22, 30, and 52 bits, respectively. These estimates combine the general dual lattice, FFT acceleration, and probability-decreasing enumeration, while following the complexity abstractions and cost models adopted in prior analyses.

Table 1. Cost estimates for the ML-KEM parameter sets. All costs are logarithms in base two.

Scheme	Attack cost		Reduction
	[21]	Ours	
ML-KEM-512	238	216	22
ML-KEM-768	347	317	30
ML-KEM-1024	478	426	52

1.2 Organization

The rest of this paper is organized as follows. Section 2 presents the necessary preliminaries, including the definition of the LWE problem, basic lattice theory, discrete Gaussian measures, and Fourier analysis tools. In Section 3 we develop our provable analysis framework in detail. We provide the equivalence characterization of the dual lattice, analyze the expectation and concentration of the score function, derive an upper bound on the false positive probability, and state the correctness theorem of the full attack. Section 4 applies the framework to the ML-KEM parameter sets, provides concrete complexity estimates, and compares with prior works. In Section 5, we compare our regime with that of [10]. Finally, we draw the conclusions in Section 6.

2 Preliminaries

Let $\mathbb{N}, \mathbb{Z}, \mathbb{R}, \mathbb{C}$ denote the sets of natural numbers, integers, real numbers, and complex numbers, respectively. Matrices are denoted by bold uppercase letters, and vectors (always column vectors) by bold lowercase letters. Unless stated otherwise, log denotes the base-2 logarithm, and $\exp(x)$ denotes the exponential function e^x. For a complex number z, $\Re(z)$ and $|z|$ denote its real part and modulus, respectively. For a finite set X, $|X|$ denotes its cardinality. $\mathbb{P}$ denotes the probability of an event. $\mathbb{E}_x[y]$ denotes the expectation of the random variable y with respect to x; when the subscript is omitted, $\mathbb{E}[y]$ denotes the expectation over all relevant random variables.

For a real number x, we define $x \bmod q$ as the unique representative in the interval $[-q/2, q/2)$ of the coset $x + q\mathbb{Z}$; namely, $x \bmod q = x - r$, where $r \in q\mathbb{Z}$ is uniquely chosen such that $x - r \in [-q/2, q/2)$. For a vector $\mathbf{x} = (x_1, \ldots, x_d)^\top$, the operation $\mathbf{x} \bmod q$ is defined component-wise. This centered reduction convention is used throughout the paper.

2.1 LWE

Definition 1 (LWE distribution). *Let $n, m, q \in \mathbb{N}$, where q denotes the modulus,and let χ_s, χ_e be distributions over $\mathbb{Z}_q^n$ and $\mathbb{Z}_q^m$ respectively. Denote by $\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}$ the probability distribution on $\mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$ obtained by sampling*

the coordinates of the matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ *independently and uniformly over* $\mathbb{Z}_q$*, sampling* $\mathbf{s} \in \mathbb{Z}_q^n$*,* $\mathbf{e} \in \mathbb{Z}_q^m$ *independently from* χ_s *and* χ_e *respectively, and outputting* $(\mathbf{A}, \mathbf{As} + \mathbf{e} \mod q)$.

The LWE problem has the following two formulations:

- Decision-LWE: Distinguish the uniform distribution over $\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ from $\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}$.
- Search-LWE: Given a sample from $\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}$, recover $\mathbf{s}$.

2.2 Discrete Fourier Transform

The Fourier transform is an operation that, for a function $f : G \to \mathbb{C}$ defined on a finite abelian group G, computes for every character χ in the dual group $\hat{G}$:

$$\hat{f}(\chi) \triangleq \sum_{g \in G} f(g)\chi(g).$$

For the specific abelian group $G = (\mathbb{Z}/q\mathbb{Z})^n$, a function $f : (\mathbb{Z}/q\mathbb{Z})^n \to \mathbb{C}$ has its Fourier transform $\hat{f}$ defined as

$$\hat{f}(v) = \sum_{u \in (\mathbb{Z}/q\mathbb{Z})^n} f(u) e^{\frac{2\pi i}{q} u^T v}.$$

Using the Fast Fourier Transform (FFT) algorithm [13], the time complexity of computing the Fourier transform for such a function is $O(nq^n)$.

2.3 Basic Concepts of Lattices

We denote by L^* the dual lattice of a lattice $L \subset \mathbb{R}^n$, defined as

$$L^* = \{x \in \mathrm{span}(L) : \forall y \in L, \langle y, x \rangle \in \mathbb{Z}\}.$$

The length of the shortest non-zero vector in L is denoted by $\lambda_1(L)$.

Let $n \in \mathbb{N}$, $1 \leq k \leq n$, and q be a prime power. A lattice L is called an n-dimensional q-ary lattice if $q\mathbb{Z}^n \subseteq L \subseteq \mathbb{Z}^n$. Given a matrix $\mathbf{A} \in \mathbb{Z}^{n \times k}$, we consider the following n-dimensional q-ary lattices:

$$L_q(\mathbf{A}) = \left\{x \in \mathbb{Z}^n : \exists s \in \mathbb{Z}^k,\, As = x \mod q\right\},$$

$$L_q^{\perp}(\mathbf{A}) = \left\{x \in \mathbb{Z}^n : A^T x = 0 \mod q\right\}.$$

For integers d and β, we denote by $BKZ_{d,\beta}$ the BKZ algorithm for a d-dimensional lattice with block size β.

2.4 Discrete Gaussian Measures on Lattices

For $\mathbf{x} \in \mathbb{R}^n$ and $s > 0$, we define the Gaussian function $\rho_s(\mathbf{x}) = \exp\left(-\frac{\pi\|\mathbf{x}\|^2}{s^2}\right)$. For a lattice $L \subset \mathbb{R}^n$, define

$$\rho_s(L) = \sum_{\mathbf{v}\in L} \rho_s(\mathbf{v}) = \sum_{\mathbf{v}\in L} \exp\left(-\frac{\pi\|\mathbf{v}\|^2}{s^2}\right).$$

The discrete Gaussian distribution $D_{L,s}$ on the lattice L with parameter $s > 0$ is defined as: for any $\mathbf{v} \in L$,

$$D_{L,s}(\mathbf{v}) = \frac{\rho_s(\mathbf{v})}{\rho_s(L)}.$$

Furthermore, we define the periodic Gaussian function $f_{L,s} : \mathbb{R}^n \to \mathbb{R}$ by

$$f_{L,s}(\mathbf{t}) = \frac{\rho_s(L+\mathbf{t})}{\rho_s(L)}.$$

Note that the value of $f_{L,s}(\mathbf{t})$ always lies in the interval $(0, 1]$.

The Poisson summation formula relates Gaussian functions on a lattice to their Fourier series on the dual lattice and is the main tool for analyzing the expectation of the score function. The following lemma states its standard form.

Lemma 1 (Poisson Summation Formula [26], Chapter VII, Proposition 15). *For an n-dimensional lattice L, let $s > 0$ and $\mathbf{t} \in \mathbb{R}^n$. Then the following holds:*

1. $\rho_s(L) = \dfrac{s^n}{\det(L)}\rho_{1/s}(L^*)$.
2. $\rho_s(L+\mathbf{t}) = \dfrac{s^n}{\det(L)} \displaystyle\sum_{\mathbf{w}\in L^*} \exp\left(2\pi i\langle \mathbf{t}, \mathbf{w}\rangle\right) \rho_{1/s}(\mathbf{w})$.

From Lemma 1, [21,24] provides a key observation: $f_{L,s}(\mathbf{t})$ is actually the expectation of a random variable defined on L^*:

$$f_{L,s}(\mathbf{t}) = \frac{\sum_{\mathbf{w}\in L^*} \exp\left(2\pi i\langle \mathbf{t}, \mathbf{w}\rangle\right) \rho_{1/s}(\mathbf{w})}{\rho_{1/s}(L^*)} = \mathbb{E}_{\mathbf{w}\leftarrow D_{L^*,1/s}}[\exp\left(2\pi i\langle \mathbf{t}, \mathbf{w}\rangle\right)].$$

Corollary 1 ([21]). *Let $L \subset \mathbb{R}^n$ be a lattice and let $s > 0$. The Fourier transform $\widehat{f_{L,s}}$ of $f_{L,s}$ is exactly the discrete Gaussian distribution $D_{L^*,1/s}$ on the dual lattice L^*. Equivalently, for any $\mathbf{w} \in L^*$, we have $\widehat{f_{L,s}}(\mathbf{w}) = \frac{\rho_{1/s}(\mathbf{w})}{\rho_{1/s}(L^*)}$.*

To bound the periodic Gaussian function $f_{L,s}(\mathbf{x})$, we express it in terms of the distance to the lattice L. The following lemma gives two inequalities for this relation.

Lemma 2 ([21], Lemma 8). *Let $L \subset \mathbb{R}^m$ be an arbitrary lattice and $s > 0$. Then for any $x \in \mathbb{R}^m$, we have*

1. *$f_{L,1/s}(x) \geq \rho_{1/s}\left(\mathrm{dist}(x, L)\right)$;*
2. *if $\mathrm{dist}(x, L) \geq \tau \triangleq \frac{1}{s}\sqrt{\frac{m}{2\pi}}$, then $f_{L,1/s}(x) \leq \rho_{1/s}\left(\mathrm{dist}(x, L) - \tau\right)$.*

2.5 Probability Inequalities

Theorem 1. *(Hoeffding's Inequality) Let* $\mathbf{w}_1, \ldots, \mathbf{w}_N$ *be independent random variables such that* $a_i \leqslant \mathbf{w}_i \leqslant b_i$ *for each* i*. Consider the sum* $S_N = \mathbf{w}_1 + \cdots + \mathbf{w}_N$*. Then for any* $t > 0$*,*

$$\Pr\left[|S_N - \mathbb{E}\left[S_N\right]| \geqslant t\right] \leqslant 2\exp\left(\frac{-2t^2}{\sum_{i=1}^{N}\left(b_i - a_i\right)^2}\right).$$

2.6 Short Vector Sampling

Exact discrete Gaussian sampling is typically impractical in lattice attack settings. We can only sample from a distribution that is statistically close to the target. The following lemma gives an efficient algorithm for such approximate sampling.

Lemma 3 (Discrete Gaussian Sampler [28]). *There is an algorithm that, given a basis* $\mathbf{B}$ *of a lattice* $L \subset \mathbb{R}^n$*, constants* $\eta > 0$*, and* $s > 0$*, outputs a sample from a distribution* $D_{L,s,\eta}$ *with* $d_{\mathrm{TV}}(D_{L,s,\eta}, D_{L,s}) \leq \eta$*. Its running time is*

$$\ln\left(\frac{1}{\eta}\right) \cdot \frac{1}{\Delta} \cdot \mathrm{poly}(n), \text{ where } \frac{1}{\Delta} = \frac{\prod_{i=1}^{n} \rho_{s/\|\bar{\mathbf{b}}_i\|}(\bar{\mathbb{Z}})}{\rho_s(L)},$$

and $\bar{\mathbf{b}}_1, \ldots, \bar{\mathbf{b}}_n$ *are the Gram-Schmidt vectors of* B*. Here* $d_{\mathrm{TV}}(D_{L,s,\eta}, D_{L,s})$ *denotes the statistical distance (or total variation distance) between* $D_{L,s,\eta}$ *and* $D_{L,s}$*.*

Black-Box 1. For integers $n \leq m$, block size β, and prime power q, there exist a deterministic algorithm B and functions T_{SV}, ℓ_{SV} such that: given $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, B runs in time $T_{\mathrm{SV}}(m, \beta, q^n)$ and returns a non-zero vector in $L_q^{\perp}(\mathbf{A})$ satisfying

$$\mathbb{E}_{\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}}\left[\|B(A)\|^2\right] \leq \ell_{\mathrm{SV}}(m, \beta, q^n)^2.$$

We sample dual vectors as follows. Given block size $2 \leq \beta \leq m$:

1. Compute a basis of $L_q^{\perp}(\mathbf{A})$;
2. Run BKZ with block size β on this basis to obtain a reduced basis $\mathbf{B}$;
3. Generate N independent samples using the MCMC Gaussian sampler [[28], Theorem 2] with parameter qs and basis B.

The complexity is

$$T_{\mathrm{sampling}}(N) = T_{\mathrm{BKZ}}(m, \beta) + N \cdot T_{\mathrm{MCMC}}\left(L_q^{\perp}(\mathbf{A}), qs\right), \tag{1}$$

where $T_{\mathrm{BKZ}}(m, \beta)$ is the cost of BKZ and $T_{\mathrm{MCMC}}(L, s)$ is the cost per sample from $D_{L,s}$. By Lemma 3,

$$T_{\mathrm{MCMC}}\left(L_q^{\perp}(\mathbf{A}), qs\right) = \ln\left(\frac{1}{\eta}\right) \cdot \frac{1}{\Delta} \cdot \mathrm{poly}(n), \quad \Delta = \frac{\rho_{qs}\left(L_q^{\perp}(\mathbf{A})\right)}{\prod_{i=1}^{n} \rho_{qs/\|\tilde{b}_i\|}(\mathbb{Z})},$$

where $\tilde{b}_1, \ldots, \tilde{b}_n$ are the Gram-Schmidt vectors of the BKZ-β reduced basis B of $L_q^{\perp}(\mathbf{A})$, and $\eta > 0$ bounds the statistical distance to the exact discrete Gaussian.

3 Our Improved Provable Dual Attack

The work of [21] represents an important step toward provable analysis of dual attacks. Subsequently, [24] introduced modulus switching into the provable dual-attack setting, enabling FFT-based evaluation over a smaller modulus, with the revised analysis given in [23]. Building on these provable frameworks and the FFT-based attack structure of [13], we extend the correctness analysis to the general dual lattice, thereby reducing the sample requirement while providing rigorous probabilistic guarantees for correctness. We first review the core analytical approach of [21], which provides the basis for the subsequent development.

3.1 The Provable Dual-Attack Framework of Pouly and Shen

For completeness and to fix notation, we recall the provable dual-attack framework of Pouly and Shen [21]. The following description summarizes the framework in the form needed for our subsequent analysis.

Let the unknown secret be $\mathbf{s} \in \mathbb{Z}_q^n$, and let $(\mathbf{A}, \mathbf{b})$ be LWE samples such that $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \mod q$, where $\mathbf{e} \in \mathbb{Z}_q^m$ is the unknown error. As in the basic dual attack, split $\mathbf{s}$ into two parts $\mathbf{s}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ and $\mathbf{s}_{\text{dual}} \in \mathbb{Z}_q^{n_{\text{dual}}}$, with $n = n_{\text{guess}} + n_{\text{dual}}$. The matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is written in block form as

$$\mathbf{A} = [\mathbf{A}_{\text{guess}} \quad \mathbf{A}_{\text{dual}}], \quad \mathbf{s} = \begin{bmatrix} \mathbf{s}_{\text{guess}} \\ \mathbf{s}_{\text{dual}} \end{bmatrix}.$$

The algorithm guesses a value $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ for $\mathbf{s}_{\text{guess}}$ and tests its correctness. Check that

$$\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}} = \mathbf{A}_{\text{guess}} \cdot (\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}}) + \mathbf{A}_{\text{dual}} \cdot \mathbf{s}_{\text{dual}} + \mathbf{e} \mod q.$$

Define the lattice

$$L_q^{\perp}(\mathbf{A}_{\text{dual}}) = \left\{ \mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \mathbf{A}_{\text{dual}} = 0 \mod q \right\}.$$

For $N \in \mathbb{N}$ and parameter $s > 0$, let $W = (\mathbf{w}_1, \ldots, \mathbf{w}_N) \in L_q^{\perp}(\mathbf{A}_{\text{dual}})^N$ be vectors sampled from $D_{L_q^{\perp}(\mathbf{A}_{\text{dual}}), qs}$. For any $\mathbf{x} \in \mathbb{R}^m$, define

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

The algorithm evaluates g_W at $\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}}$ for every $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ and outputs the guess that maximizes this value. The analysis proceeds as follows.

Set $L = L_q(\mathbf{A}_{\text{dual}})$ to simplify notation. Recall the periodic Gaussian function from Section 2.4:

$$f_{L,1/s}(\mathbf{x}) = \frac{\rho_{1/s}(\mathbf{x} + L)}{\rho_{1/s}(L)},$$

where $\mathbf{x} \in \mathbb{R}^m$ and $s > 0$. For sufficiently large N, g_W approximates $f_{L,1/s}$ on integer vectors with high probability. Since $\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}}$ is an integer vector, the behavior of $f_{L,1/s}$ determines the output. Standard Gaussian tail bounds give the approximation

$$f_{L,1/s}(\mathbf{x}) \approx \rho_{1/s}(\text{dist}(\mathbf{x}, L))$$

for any $s > 0$ and $\mathbf{x} \in \mathbb{R}^m$. Thus $f_{L,1/s}(\mathbf{x})$ measures the distance from $\mathbf{x}$ to L.

The attack relies on the following distance inequality: for almost all choices of $\mathbf{A}$ and $\mathbf{e}$, and for any $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}} \setminus \{\mathbf{s}_{\text{guess}}\}$,

$$\text{dist}(\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \mathbf{s}_{\text{guess}}, L) < \text{dist}(\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}}, L). \tag{2}$$

This implies

$$f_{L,1/s}(b - \mathbf{A}_{\text{guess}} \cdot \mathbf{s}_{\text{guess}}) > f_{L,1/s}(b - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}}),$$

and for large N the function g_W preserves this ordering, so the algorithm recovers $\mathbf{s}_{\text{guess}}$.

To formalize inequality (2), a separation condition is required. The work of [21] shows: if $2\|\mathbf{e}\| < \lambda_1(L_q(\mathbf{A}))$, then for every $\mathbf{x} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L$,

$$f_{L,1/s}(\mathbf{e}) > f_{L,1/s}(\mathbf{e} + \mathbf{x}). \tag{3}$$

The proof of inequality (3) is as follows. Since $\mathbf{A}_{\text{dual}} \cdot \mathbf{s}_{\text{dual}} \in L$, the left side of inequality (2) expands to

$$\text{dist}(\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \mathbf{s}_{\text{guess}}, L) = \text{dist}(\mathbf{e} + \mathbf{A}_{\text{dual}} \cdot \mathbf{s}_{\text{dual}}, L) = \text{dist}(\mathbf{e}, L) \leq \|\mathbf{e}\|.$$

For an incorrect guess $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$,

$$\text{dist}(\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}}, L) = \text{dist}(\mathbf{e} + \mathbf{A}_{\text{dual}} \cdot \mathbf{s}_{\text{dual}} + \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}}), L).$$

This simplifies to $\text{dist}(\mathbf{e} + \mathbf{x}, L)$ with $\mathbf{x} = \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}}) \in L_q(\mathbf{A}_{\text{guess}})$. Assume $\mathbf{x} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L$. Then

$$\begin{aligned}
\text{dist}(\mathbf{e} + \mathbf{x}, L) &= \min\{\|\mathbf{e} + \mathbf{x} + \mathbf{z}\| : \mathbf{z} \in L\} \\
&\geq \min\{\|\mathbf{e} + \mathbf{y} + \mathbf{z}\| : \mathbf{z} \in L, \mathbf{y} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L\} \\
&\geq \min\{\|\mathbf{y} + \mathbf{z}\| : \mathbf{z} \in L, \mathbf{y} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L\} - \|\mathbf{e}\| \\
&\geq \lambda_1\left(L + L_q(\mathbf{A}_{\text{guess}})\right) - \|\mathbf{e}\|.
\end{aligned}$$

The last step uses that $\mathbf{y} + \mathbf{z} \neq 0$ for $\mathbf{z} \in L$ and $\mathbf{y} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L$. Inequality (2) therefore holds whenever

$$\|\mathbf{e}\| < \lambda_1\left(L + L_q(\mathbf{A}_{\text{guess}})\right) - \|\mathbf{e}\|.$$

Since $L + L_q(\mathbf{A}_{\text{guess}}) = L_q(\mathbf{A}_{\text{dual}}) + L_q(\mathbf{A}_{\text{guess}}) = L_q(\mathbf{A})$, this condition is exactly $2\|\mathbf{e}\| < \lambda_1(L_q(\mathbf{A}))$, which yields inequality (3).

To apply the result, note that inequality (3) requires $\mathbf{x} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L$. Here $\mathbf{x} = \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}})$ with $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$. The condition $\mathbf{x} \notin L$ is equivalent to $\mathbf{x} \notin \mathbf{A}_{\text{dual}}\mathbb{Z}_q^{n_{\text{dual}}} + q\mathbb{Z}^m$. If $\mathbf{A}$ has full rank over $\mathbb{Z}_q$ (an event of overwhelming probability), this holds.

Thus the dual attack of [21] succeeds for essentially all instances $(\mathbf{A}, \mathbf{e})$ satisfying $2\|\mathbf{e}\| < \lambda_1(L_q(\mathbf{A}))$. This condition ensures that $\mathbf{e}$ is the unique shortest non-zero vector in $L_q(\mathbf{A})$, which underlies the perfect-recovery argument.

3.2 Our Provable Dual Attack

We adopt the general dual-attack structure of [13], which partitions the secret into three parts and uses FFT to evaluate the score function over a large guessing space. We incorporate this structure into the provable correctness framework and adapt it to the general dual lattice considered below.

A central motivation for working with this structure is its sample efficiency. Prior work on provable dual attacks [21,23] is restricted to the orthogonal lattice $L_q^{\perp}(\mathbf{A}) = \{\mathbf{w} \in \mathbb{Z}_q^m \mid \mathbf{w}^{\mathrm{T}}\mathbf{A} = \mathbf{0},\ \mathbf{A} \in \mathbb{Z}_q^{m\times n}\}$, which requires approximately $m \approx 2n$ LWE samples to achieve optimal time complexity. In many practical scenarios, however, only n samples are available. By working instead with the general dual lattice

$$L_q(\mathbf{A}^T) = \{(\mathbf{w}, \mathbf{y}) \in \mathbb{Z}_q^{m+n} \mid \mathbf{w}^{\mathrm{T}}\mathbf{A} = \mathbf{y}^{\mathrm{T}} \mod q\},$$

from which we sample short vectors, the sample requirement drops to $m \approx n$. We now detail how we extend the correctness analysis to this lattice.

Let $\mathbf{s} \in \mathbb{Z}_q^n$ be the unknown secret, and let $(\mathbf{A}, \mathbf{b})$ denote a set of LWE samples with $\mathbf{b} = \mathbf{A}\mathbf{s}+\mathbf{e} \mod q$, where $\mathbf{e} \in \mathbb{Z}_q^m$ is the unknown error. Partition the secret $\mathbf{s}$ into three parts $\mathbf{s}_0 \in \mathbb{Z}_q^{n_0}$, $\mathbf{s}_1 \in \mathbb{Z}_q^{n_1}$ and $\mathbf{s}_2 \in \mathbb{Z}_q^{n_2}$, where $n = n_0 + n_1 + n_2$. The matrix $\mathbf{A} \in \mathbb{Z}_q^{m\times n}$ is partitioned accordingly as

$$\mathbf{A} = \begin{bmatrix}\mathbf{A}_0 & \mathbf{A}_1 & \mathbf{A}_2\end{bmatrix}, \quad \mathbf{s} = \begin{bmatrix}\mathbf{s}_0 \\ \mathbf{s}_1 \\ \mathbf{s}_2\end{bmatrix}.$$

The dual attack is used only to recover $\mathbf{s}_0$. Following [13], let γ be a small element of the ring $\mathbb{Z}_q$ such that $\gamma \cdot \rho = \pm 1 \mod q$ for some element ρ; hence γ^{-1} is well defined and equals either ρ or $-\rho$. In practice, the modulus q is typically chosen to be a prime or a power of two. When q is odd prime we may take $\gamma = 2$; for the power-of-two case we take $\gamma = 3$. We illustrate the operation of the distinguisher in the prime modulus case. Set $\gamma = 2$ and define $\hat{\mathbf{A}}_1 = \gamma\mathbf{A}_1 \mod q$. Consider the dual lattice

$$L_q(\mathbf{A}_{\mathrm{dual}}^T) = \{(\mathbf{x}, \mathbf{y}_1, \mathbf{y}_2) \in \mathbb{Z}^m \times \mathbb{Z}^{n_1} \times \mathbb{Z}^{n_2} : (\hat{\mathbf{A}}_1 \mid \mathbf{A}_2)^T\mathbf{x} = \begin{pmatrix}\mathbf{y}_1 \\ \mathbf{y}_2\end{pmatrix} \mod q\}.$$

This choice of the dual lattice differs slightly from that of [13]; we adopt this variant because it is more convenient for the provable analysis that follows. The dimension of this lattice is $d = m + n_1 + n_2 = m + n - n_0$.

The Dual Attack Procedure We now describe the dual attack itself. For notational convenience, set $\theta = \exp(2\pi i/q)$ and define the following inner products:

$$b_j = \mathbf{w}_j^T\mathbf{b}, \quad e_j = \mathbf{w}_j^T\mathbf{e}, \quad \mathbf{v}_{0,j} = \mathbf{w}_j^T\mathbf{A}_0, \quad \hat{\mathbf{v}}_{1,j} = \mathbf{w}_j^T\hat{\mathbf{A}}_1, \quad \mathbf{v}_{2,j} = \mathbf{w}_j^T\mathbf{A}_2.$$

For each guess $(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1)$ of $(\mathbf{s}_0, \mathbf{s}_1)$, the score function is defined as

$$F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) = \frac{1}{N}\sum_{j=1}^{N} \cos\left(\frac{2\pi}{q}\big(b_j - \mathbf{v}_{0,j}^T\tilde{\mathbf{s}}_0 - \hat{\mathbf{v}}_{1,j}^T\tilde{\mathbf{s}}_1 \cdot q/2\big)\right)$$

$$= \frac{1}{N}\Re\left(\sum_{j=1}^{N} \theta^{\,b_j - \mathbf{v}_{0,j}^T\tilde{\mathbf{s}}_0 - \hat{\mathbf{v}}_{1,j}^T\tilde{\mathbf{s}}_1 \cdot q/2}\right) = \frac{1}{N}\Re\left(\sum_{j=1}^{N} \theta^{\,b_j - \mathbf{v}_{0,j}^T\tilde{\mathbf{s}}_0} \cdot \exp\big(-\hat{\mathbf{v}}_{1,j}^T\tilde{\mathbf{s}}_1 \cdot \pi i\big)\right).$$

Since $\exp(\pi i \cdot \hat{\mathbf{v}}_{1,j}^T\tilde{\mathbf{s}}_1) = (-1)^{\langle \hat{\mathbf{v}}_{1,j} \mod 2, \tilde{\mathbf{s}}_1\rangle}$, the samples can be grouped according to the parity vector $\mathbf{c} = \hat{\mathbf{v}}_{1,j} \mod 2 \in \mathbb{Z}_2^{n_1}$. Define

$$f_{\mathbf{W},\tilde{\mathbf{s}}_0}(\mathbf{c}) = \sum_{j \in I(\mathbf{c})} \theta^{\,b_j - \mathbf{v}_{0,j}^T\tilde{\mathbf{s}}_0},$$

where $I(\mathbf{c}) = \{j \mid \hat{\mathbf{v}}_{1,j} \mod 2 = \mathbf{c}\}$. The score function can then be rewritten as

$$F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) = \frac{1}{N}\Re\left(\sum_{\mathbf{c} \in \mathbb{Z}_2^{n_1}} f_{\mathbf{W},\tilde{\mathbf{s}}_0}(\mathbf{c})\,(-1)^{-\langle \mathbf{c}, \tilde{\mathbf{s}}_1\rangle}\right).$$

Let $\mathbb{S}_0$ denote the search space for $\mathbf{s}_0$. For each fixed $\tilde{\mathbf{s}}_0 \in \mathbb{S}_0$, the attacker exhaustively tries all binary vectors $\tilde{\mathbf{s}}_1 \in \mathbb{Z}_2^{n_1}$ and evaluates the corresponding values $F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1)$. This can be performed in time $O(N + 2^{n_1})$ by means of the fast Fourier transform. The algorithm then outputs

$$(\mathbf{s}_0, \mathbf{s}_1) = \underset{\tilde{\mathbf{s}}_0 \in \mathbb{S}_0,\, \tilde{\mathbf{s}}_1 \in \mathbb{Z}_2^{n_1}}{\operatorname{argmax}} \big(F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1)\big),$$

With sufficiently many samples, the maximizer coincides with $(\mathbf{s}_0, \mathbf{s}_1 \mod 2)$, although we only claim the recovery of $\mathbf{s}_0$.

At this stage, the adversary has recovered $\mathbf{s}_0$. Write the secret as $\mathbf{s} = \mathbf{s}_0 \,\|\, \mathbf{s}'$, where $\mathbf{s}' \in \mathbb{Z}_q^{n-n_0}$ is the remaining part. Partition $\mathbf{A}$ accordingly as $\mathbf{A} = [\mathbf{A}_0 \mid \mathbf{A}']$, so that

$$\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} = \mathbf{A}_0\mathbf{s}_0 + \mathbf{A}'\mathbf{s}' + \mathbf{e} \mod q.$$

Subtracting $\mathbf{A}_0\mathbf{s}_0$ yields a new LWE instance

$$\mathbf{b}' = \mathbf{b} - \mathbf{A}_0\mathbf{s}_0 = \mathbf{A}'\mathbf{s}' + \mathbf{e} \mod q,$$

of dimension $(m, n-n_0)$ with secret $\mathbf{s}'$. For the parameter sets considered in this paper, a sufficient number of bits are recovered in the first call, so the remaining secret is of much lower dimension and the cost of further calls to the dual attack becomes negligible.

Correctness Analysis To analyze the correctness of the dual attack from [13] using the framework of [21], it is necessary to transform the dual lattice employed in the attack into an orthogonal lattice that is more amenable to analysis. We establish the following equivalence. The lemma shows that the dual lattice $L_q(\mathbf{A}_{\text{dual}}^T)$ and a certain orthogonal lattice $L_q^{\perp}(\mathbf{B}_{\text{dual}})$ represent the same q-ary lattice, which allows us to exploit known properties of orthogonal lattices.

Lemma 4. *As subsets of* $\mathbb{Z}^m \times \mathbb{Z}^{n_1} \times \mathbb{Z}^{n_2}$*, we have* $L_q(\mathbf{A}_{dual}^T) = L_q^\perp(\mathbf{B}_{dual})$*, where*

$$L_q(\mathbf{A}_{dual}^T) = \left\{ (\mathbf{x}, \mathbf{y}_1, \mathbf{y}_2) \in \mathbb{Z}^m \times \mathbb{Z}^{n_1} \times \mathbb{Z}^{n_2} \;\middle|\; (\hat{\mathbf{A}}_1 \mid \mathbf{A}_2)^T \mathbf{x} \equiv \begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} \mod q \right\},$$

and

$$L_q^\perp(\mathbf{B}_{dual}) = \left\{ \mathbf{z} \in \mathbb{Z}^{m+n-n_0} \;\middle|\; \mathbf{B}_{dual}^T \mathbf{z} \equiv 0 \mod q \right\}, \quad \mathbf{B}_{dual} = \begin{pmatrix} \hat{\mathbf{A}}_1 & \mathbf{A}_2 \\ -\mathbf{I}_{n_1} & 0 \\ 0 & -\mathbf{I}_{n_2} \end{pmatrix}.$$

Proof. See Appendix A. □

This equivalence is a key technical insight: it allows us to transplant the provable framework of [21], originally developed for orthogonal lattices, to the more general and sample-efficient dual lattice $L_q(\mathbf{A}_{\text{dual}}^T)$.

As in standard dual attacks, we employ a short-vector sampling algorithm to obtain independent short vectors $\mathbf{w}_j \in L_q(\mathbf{A}_{\text{dual}}^T)$ for $j \in [N]$. By Lemma 3, these induce vectors

$$\mathbf{v}_j \triangleq \begin{pmatrix} \mathbf{w}_j \\ \hat{\mathbf{A}}_1^T \mathbf{w}_j \\ \mathbf{A}_2^T \mathbf{w}_j \end{pmatrix} \mod q \sim D_{L_q^\perp(\mathbf{B}_{\text{dual}}), qs}, \quad j = 1, \ldots, N.$$

For the correctness analysis, it will be convenient to work with $\hat{\mathbf{s}}_1 = \gamma^{-1}\mathbf{s}_1 = \frac{q+1}{2}\mathbf{s}_1 \mod q$. In the prime modulus case ($\gamma = 2$), this simplifies to $\hat{\mathbf{s}}_1 = \frac{q+1}{2}\mathbf{s}_1 \mod q$. From the LWE relation we have, over the integers,

$$b_j = \mathbf{v}_{0,j}^T \mathbf{s}_0 + \hat{\mathbf{v}}_{1,j}^T \hat{\mathbf{s}}_1 + \mathbf{v}_{2,j}^T \mathbf{s}_2 + e_j + \ell_j q \quad \text{for some } \ell_j \in \mathbb{Z},$$

Substituting the definition of $\hat{\mathbf{s}}_1$ gives

$$b_j = \mathbf{v}_{0,j}^T \mathbf{s}_0 + \hat{\mathbf{v}}_{1,j}^T \cdot \tfrac{q+1}{2}\, \mathbf{s}_1 + \mathbf{v}_{2,j}^T \mathbf{s}_2 + e_j + \ell_j q.$$

For a guess $(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1)$ we obtain

$$\begin{aligned} & b_j - \mathbf{v}_{0,j}^T \tilde{\mathbf{s}}_0 - \hat{\mathbf{v}}_{1,j}^T \tilde{\mathbf{s}}_1 \cdot \frac{q}{2} \\ =& \mathbf{v}_{0,j}^T (\mathbf{s}_0 - \tilde{\mathbf{s}}_0) + \hat{\mathbf{v}}_{1,j}^T \Big(\frac{q+1}{2} \mathbf{s}_1 - \frac{q}{2} \tilde{\mathbf{s}}_1 \Big) + \mathbf{v}_{2,j}^T \mathbf{s}_2 + e_j + \ell_j q \\ =& \mathbf{v}_{2,j}^T \mathbf{s}_2 + \frac{1}{2} \hat{\mathbf{v}}_{1,j}^T \mathbf{s}_1 + e_j + \frac{q}{2} \hat{\mathbf{v}}_{1,j}^T (\mathbf{s}_1 - \tilde{\mathbf{s}}_1) + \mathbf{v}_{0,j}^T (\mathbf{s}_0 - \tilde{\mathbf{s}}_0) + \ell_j q. \end{aligned}$$

Substituting this expansion into the cosine form of the score function, we obtain, for $\tilde{\mathbf{s}}_0 \in \mathbb{S}_0$ and $\tilde{\mathbf{s}}_1 \in \mathbb{Z}_2^{n_1}$, the score function can be written as follows:

$$
\begin{aligned}
F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) &= \frac{1}{N}\sum_{j=1}^{N} \cos\left(\frac{2\pi}{q}\left\langle \mathbf{b} - \mathbf{A}_0\tilde{\mathbf{s}}_0 - \hat{\mathbf{A}}_1\tilde{\mathbf{s}}_1\frac{q}{2},\, \mathbf{w}_j\right\rangle\right) \\
&= \frac{1}{N}\sum_{j=1}^{N} \cos\left(\frac{2\pi}{q}\left\langle \mathbf{A}_2\mathbf{s}_2 + \frac{1}{2}\hat{\mathbf{A}}_1\mathbf{s}_1 + \mathbf{e} + \frac{q}{2}\hat{\mathbf{A}}_1(\mathbf{s}_1 - \tilde{\mathbf{s}}_1) + \mathbf{A}_0(\mathbf{s}_0 - \tilde{\mathbf{s}}_0),\, \mathbf{w}_j\right\rangle\right) \\
&= \frac{1}{N}\sum_{j=1}^{N} \cos\left(\frac{2\pi}{q}\left\langle \begin{pmatrix}\mathbf{e}\\ \mathbf{s}_1/2\\ \mathbf{s}_2\end{pmatrix} + \begin{pmatrix}\mathbf{A}_0(\mathbf{s}_0 - \tilde{\mathbf{s}}_0)\\ \frac{q}{2}(\mathbf{s}_1 - \tilde{\mathbf{s}}_1)\\ \mathbf{0}\end{pmatrix},\, \begin{pmatrix}\mathbf{w}_j\\ \hat{\mathbf{v}}_{1,j}\\ \mathbf{v}_{2,j}\end{pmatrix}\right\rangle\right) \\
&\triangleq \frac{1}{N}\sum_{j=1}^{N} \cos\left(\frac{2\pi}{q}\left\langle \mathbf{e}_{\mathrm{lat}} + \Delta_{\mathrm{lat}},\, \mathbf{v}_j\right\rangle\right).
\end{aligned}
$$

Here $\mathbf{v}_1, \ldots, \mathbf{v}_N \in L_q^{\perp}(\mathbf{B}_{\mathrm{dual}})$ are independent and identically distributed samples from the discrete Gaussian distribution $D^N_{L_q^{\perp}(\mathbf{B}_{\mathrm{dual}}),qs}$. The score function can therefore be rewritten as

$$
F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) = \frac{1}{N}\sum_{j=1}^{N} \cos\left(2\pi\left\langle \mathbf{t},\, \frac{1}{q}\mathbf{v}_j\right\rangle\right),
$$

where

$$
\mathbf{t} = \mathbf{e}_{\mathrm{lat}} + \Delta_{\mathrm{lat}}, \quad \mathbf{e}_{\mathrm{lat}} = \begin{pmatrix}\mathbf{e}\\ \frac{1}{2}\mathbf{s}_1\\ \mathbf{s}_2\end{pmatrix}, \quad \Delta_{\mathrm{lat}} = \begin{pmatrix}\mathbf{A}_0(\mathbf{s}_0 - \tilde{\mathbf{s}}_0)\\ \frac{q}{2}(\mathbf{s}_1 - \tilde{\mathbf{s}}_1)\\ \mathbf{0}\end{pmatrix}.
$$

Thus, we have successfully reformulated the heuristic FFT-based score function as the empirical average of cosines of inner products with i.i.d. short vectors drawn from an orthogonal lattice. This expression serves as the starting point for our rigorous probabilistic analysis.

A natural starting point for our correctness analysis, given the lattice equivalence established in Lemma 4, would be to directly apply the geometric framework of [21]. In their work, the attack's success is guaranteed by a clean distance-based condition: if the noise vector $\mathbf{e}$ satisfies $2\|\mathbf{e}\| < \lambda_1(L_q(\mathbf{A}))$, then the score of the correct guess will provably exceed that of all incorrect ones. This elegant criterion links the correctness directly to a well-understood property of the *uniformly random* q-ary lattice $L_q(\mathbf{A})$, whose first minimum can be sharply estimated via the Gaussian Heuristic.

However, a key difficulty in our setting is that the distinguishing lattice $L_q(\mathbf{B}_{\mathrm{dual}})$ is no longer a uniformly random q-ary lattice. Indeed, due to the fixed identity block and the scaled submatrix appearing in $\mathbf{B}_{\mathrm{dual}}$, the Gaussian heuristic cannot be directly applied to estimate $\lambda_1(L_q(\mathbf{B}_{\mathrm{dual}}))$. Consequently, rather than attempting to characterize the exact first minimum of the lattice,

we formulate the correctness condition in terms of a geometric separation margin.

Motivated by this difficulty, we seek a multiplicative separation condition ensuring that every incorrect target point remains significantly farther from the lattice than the correct one. Concretely, we shall require a large enough constant $c > 1$ such that the correct target point is separated from every incorrect target point by a multiplicative distance gap:

$$\operatorname{dist}\left(\mathbf{e}_{\mathrm{lat}} + \Delta_{\mathrm{lat}}, L_q(\mathbf{B}_{\mathrm{dual}})\right) > c \cdot \operatorname{dist}\left(\mathbf{e}_{\mathrm{lat}}, L_q(\mathbf{B}_{\mathrm{dual}})\right)$$

The parameter c therefore quantifies the geometric distinguishability between correct and incorrect guesses in the lattice.

Under such a separation condition, the periodic Gaussian values associated with incorrect guesses become sufficiently smaller than those of the correct guess, enabling reliable distinction via the score function. The remainder of this subsection is devoted to formalizing this separation property and deriving the corresponding correctness and false-positive bounds.

Now, having expressed the score as an average of bounded random variables, two properties govern its behavior: its expectation and its concentration. Once the score function is written as an average of cosines, the next step is to expand its expectation via a Fourier series. This requires the duality relation between $L_q(\mathbf{B}_{\mathrm{dual}})$ and $L_q^{\perp}(\mathbf{B}_{\mathrm{dual}})$, which is given by the following lemma.

Lemma 5. *Let q be a prime, and set $M = m + n_1 + n_2$, $N = n_1 + n_2$. Define the q-ary lattices*

$$L_q(\mathbf{B}_{\mathrm{dual}}) = \left\{\mathbf{B}_{\mathrm{dual}}\mathbf{x} + q\mathbf{y} \in \mathbb{R}^M : \mathbf{x} \in \mathbb{Z}^N,\ \mathbf{y} \in \mathbb{Z}^M\right\},$$

$$L_q^{\perp}(\mathbf{B}_{\mathrm{dual}}) = \left\{\mathbf{y} \in \mathbb{Z}^M : \mathbf{B}_{\mathrm{dual}}^T\mathbf{y} \in q\mathbb{Z}^N\right\}.$$

Then the following duality relations hold:

$$\left(L_q(\mathbf{B}_{\mathrm{dual}})\right)^* = \frac{1}{q}L_q^{\perp}(\mathbf{B}_{\mathrm{dual}}), \qquad \left(L_q^{\perp}(\mathbf{B}_{\mathrm{dual}})\right)^* = \frac{1}{q}L_q(\mathbf{B}_{\mathrm{dual}}).$$

By Lemma 5, we have $(L_q(\mathbf{B}_{\mathrm{dual}}))^* = \frac{1}{q}L_q^{\perp}(\mathbf{B}_{\mathrm{dual}})$. Since $\mathbf{v}_j \sim D_{L_q^{\perp}(\mathbf{B}_{\mathrm{dual}}),qs}$, it follows that $\frac{1}{q}\mathbf{v}_j \sim D_{(L_q(\mathbf{B}_{\mathrm{dual}}))^*,s}$. Applying the Fourier expansion of Corollary 1 to the periodic Gaussian function $f_{L_q(\mathbf{B}_{\mathrm{dual}}),1/s}$ yields the following lemma.

Lemma 6 (Fourier Series). *For any $\mathbf{t} \in \mathbb{R}^M$,*

$$f_{L_q(\mathbf{B}_{\mathrm{dual}}),1/s}(\mathbf{t}) = \sum_{\mathbf{w} \in (L_q(\mathbf{B}_{\mathrm{dual}}))^*} D_{(L_q(\mathbf{B}_{\mathrm{dual}}))^*,s}(\mathbf{w}) \cos(2\pi\langle \mathbf{t}, \mathbf{w}\rangle).$$

By Lemma 6 and the definition of the score function as an average of cosines, we immediately obtain, for any fixed $\mathbf{t}$, $\mathbb{E}[F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1)] = f_{L_q(\mathbf{B}_{\mathrm{dual}}),1/s}(\mathbf{t})$.

Lemma 6 identifies the expectation of the score function with the periodic Gaussian function $f_{L_q(\mathbf{B}_{\mathrm{dual}}),1/s}(\mathbf{t})$. It therefore remains to control the concentration of the empirical score around this expectation. The following lemma shows that, for sufficiently large N, the approximation holds uniformly over the entire search space.

Lemma 7 (Approximation of the Score Function Using the Periodic Gaussian Function). *Let $s > 0$, $\delta > 0$, and $N \in \mathbb{N}$. Define the set of target points as*

$$X_{\text{total}} = \left\{ \mathbf{e}_{\text{lat}} + \Delta_{\text{lat}} \mid \tilde{\mathbf{s}}_0 \in \mathbb{S}_0,\ \tilde{\mathbf{s}}_1 \in \mathbb{Z}_2^{n_1} \right\}$$
$$= \left\{ \begin{pmatrix} \mathbf{e} \\ \mathbf{s}_1/2 \\ \mathbf{s}_2 \end{pmatrix} + \begin{pmatrix} \mathbf{A}_0^T \cdot (\mathbf{s}_0 - \tilde{\mathbf{s}}_0) \\ \frac{q}{2} \cdot (\mathbf{s}_1 - \tilde{\mathbf{s}}_1) \\ \mathbf{0} \end{pmatrix} \;\middle|\; \tilde{\mathbf{s}}_0 \in \mathbb{S}_0,\ \tilde{\mathbf{s}}_1 \in \mathbb{Z}_2^{n_1} \right\},$$

whose cardinality is $|X_{\text{total}}| = |\mathbb{S}_0| \cdot 2^{n_1}$. Let $\mathbf{v}_1, \ldots, \mathbf{v}_N \in L_q^{\perp}(\mathbf{B}_{\text{dual}})$ be independent samples from $D^N_{L_q^{\perp}(\mathbf{B}_{\text{dual}}),qs}$. Denote

$$\mathcal{E} := \max_{(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) \in \mathbb{S}_0 \times \mathbb{Z}_2^{n_1}} \left| F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) - f_{L_q(\mathbf{B}_{\text{dual}}),1/s}(\mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}) \right|.$$

Then $\Pr(\mathcal{E} \le \delta) \ge 1 - |\mathbb{S}_0| \cdot 2^{n_1} \cdot 2^{-\Omega(N\delta^2)}$.

Proof. From Lemma 6 we have $\mathbb{E}[F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1)] = f_{L_q(\mathbf{B}_{\text{dual}}),1/s}(\mathbf{t})$. To analyze the concentration of the score function around its expectation, fix an arbitrary target point $\mathbf{t} \in X_{\text{total}}$. Write the score function as $F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) = \frac{1}{N} \sum_{j=1}^{N} h_{\mathbf{v}_j}(\mathbf{t})$, where $h_{\mathbf{v}_j}(\mathbf{t}) = \cos(2\pi \langle \mathbf{t}, \frac{1}{q}\mathbf{v}_j \rangle)$. Since the cosine function takes values in $[-1, 1]$, we have $h_{\mathbf{v}_j}(\mathbf{t}) \in [-1, 1]$ for all j. Applying theorem 1 to the sum of N independent random variables bounded in $[-1, 1]$ yields, for any fixed $\mathbf{t}$,

$$\Pr\left(\left| \frac{1}{N} \sum_{j=1}^{N} h_{\mathbf{v}_j}(\mathbf{t}) - \mathbb{E}[h_{\mathbf{v}_1}(\mathbf{t})] \right| \ge \delta \right) \le 2 \exp\left(-\frac{N\delta^2}{2} \right) = 2^{-\Omega(N\delta^2)}.$$

We now consider all possible guesses simultaneously. The set X_{total} contains $|\mathbb{S}_0| \cdot 2^{n_1}$ points. By the union bound,

$$\Pr\left(\exists \mathbf{t} \in X_{\text{total}},\ \left| F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) - f_{L_q(\mathbf{B}_{\text{dual}}),1/s}(\mathbf{t}) \right| \ge \delta \right) \le |\mathbb{S}_0| \cdot 2^{n_1} \cdot 2^{-\Omega(N\delta^2)}.$$

Taking the complement yields the claimed bound. □

Lemma 7 shows that, with sufficiently many short vectors, the score function uniformly approximates the periodic Gaussian function $f_{L_q(\mathbf{B}_{\text{dual}}),1/s}$ on the entire set of target points. Hence, to guarantee that the correct guess attains a strictly higher score than all incorrect guesses, it suffices to compare the values of $f_{L_q(\mathbf{B}_{\text{dual}}),1/s}$ at the corresponding lattice points.

From Lemma 7 and Lemma 2 we obtain the following corollary, which translates the comparison of score functions into distance bounds and will serve as the key tool for bounding false positives.

Corollary 2. *Let $\tau = \frac{1}{s}\sqrt{\frac{m+n-n_0}{2\pi}}$. Then for all $(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) \in \mathbb{S}_0 \times \mathbb{Z}_2^{n_1}$ and with $\mathbf{t} = \mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}$, the following hold:*

1. $f_{L_q(\mathbf{B}_{\text{dual}}),1/s}(\mathbf{e}_{\text{lat}}) \geq \rho_{1/s}\big(\text{dist}(\mathbf{e}_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}}))\big)$;
2. *If* $\text{dist}(\mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}})) \geq \tau$, *then*

$$f_{L_q(\mathbf{B}_{\text{dual}}),1/s}(\mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}) \leq \rho_{1/s}\big(\text{dist}(\mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}, L) - \tau\big).$$

Corollary 2 reduces the correctness analysis to a geometric separation problem: the correct target point must remain significantly closer to the lattice than any incorrect one. We next formalize this separation condition and derive corresponding bounds on the false positive probability.

Observe that the function $\rho_{1/s}$ is monotonically decreasing in its argument. Hence, to guarantee that the algorithm successfully distinguishes the correct guess from all incorrect ones, it suffices to have a constant $c > 1$ such that for every $\Delta_{\text{lat}} \neq \mathbf{0}$,

$$c \cdot \text{dist}\big(\mathbf{e}_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}})\big) < \text{dist}\big(\mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}})\big). \tag{4}$$

Moreover, the constant c must satisfy $c \cdot \text{dist}(\mathbf{e}_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}})) \geq \tau + 3\|\mathbf{e}_{\text{lat}}\|$. Indeed, since the accuracy parameter δ is typically chosen to be much smaller than $\rho_{1/s}\big(\text{dist}(\mathbf{e}_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}}))\big)$, when c meets the above requirement and the concentration bound of Lemma 7 holds. This event occurs with sufficiently high probability when N is sufficiently large. The following chain of inequalities holds:

$$\begin{aligned} F_{\mathbf{W}}(\mathbf{s}_0, \mathbf{s}_1) \geq & f_{L_q(\mathbf{B}_{\text{dual}}),1/s}(\mathbf{e}_{\text{lat}}) - \delta \\ \geq & \rho_{1/s}\big(\text{dist}(\mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}, L) - \tau\big) + \delta \geq F_{\mathbf{W}}(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1), \end{aligned}$$

which implies that the score of the correct guess strictly exceeds that of any incorrect guess.

If Condition (4) is violated for some incorrect guess $(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1)$, then the geometric separation required for correctness breaks down. In that case, a false positive may occur. We therefore define the geometric separation event

$$\mathcal{G}: \quad \forall \Delta_{\text{lat}} \neq \mathbf{0}, \text{dist}\left(\mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}})\right) > c \cdot \text{dist}\left(\mathbf{e}_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}})\right)$$

Under this event, the correct guess achieves a strictly larger score than every incorrect guess with sufficiently high probability. Accordingly, a false positive can occur only if the event $\mathcal{G}$ fails.

To assess the failure probability of event $\mathcal{G}$, we relax the distance inequality into a more tractable algebraic form. For the right-hand side, we use the trivial upper bound $\text{dist}(\mathbf{e}_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}})) \leq \|\mathbf{e}_{\text{lat}}\|$. For the left-hand side, we lower-bound the distance by the minimum over an explicit set of integer vectors:

$$\begin{aligned} \text{dist}\big(\mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}, L_q(\mathbf{B}_{\text{dual}})\big) &= \min_{\substack{\tilde{\mathbf{s}}_0 \in \mathbb{S}_0,\, \tilde{\mathbf{s}}_1 \in \mathbb{Z}_2^{n_1}, \\ \mathbf{z} \in L_q(\mathbf{B}_{\text{dual}})}} \|\mathbf{e}_{\text{lat}} + \Delta_{\text{lat}} - \mathbf{z}\| \\ &\geq \min_{\substack{\mathbf{x}_i \in \mathbb{Z}^{n_i}, \\ \mathbf{y}_1 \in \mathbb{Z}^{n_1}}} \left\| \begin{pmatrix} \mathbf{A}_0\mathbf{x}_0 + \hat{\mathbf{A}}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2 \\ \frac{q}{2}\mathbf{y}_1 - \mathbf{x}_1 \\ -\mathbf{x}_2 \end{pmatrix} + \begin{pmatrix} \mathbf{e} \\ \mathbf{s}_1/2 \\ \mathbf{s}_2 \end{pmatrix} \mod q \right\|. \end{aligned}$$

Equivalently, the event $\mathcal{G}$ fails if there exists an incorrect guess $(\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1) \neq (\mathbf{s}_0, \mathbf{s}_1 \mod 2)$ such that the corresponding score is at least that of the correct guess. By Lemma 7, this situation is equivalent to the existence of non-zero vectors $\Delta_0 = \mathbf{s}_0 - \tilde{\mathbf{s}}_0 \in \mathbb{Z}_q^{n_0}$ and $\Delta_1 = \mathbf{s}_1 - \tilde{\mathbf{s}}_1 \in \mathbb{Z}_2^{n_1}$ such that the distance from the point $\mathbf{t}_\Delta = \mathbf{e}_{\text{lat}} + \Delta_{\text{lat}}$ to the lattice $L_q(\mathbf{B}_{\text{dual}})$ does not satisfy the separation condition (4). The following lemma derives a necessary condition for the failure of $\mathcal{G}$.

Lemma 8 (Necessary Condition for Geometric Separation Failure). *A necessary condition for a false positive event is the existence of integer vectors $\mathbf{x}_0 \in \mathbb{Z}^{n_0}$, $\mathbf{x}_1 \in \mathbb{Z}^{n_1}$, $\mathbf{x}_2 \in \mathbb{Z}^{n_2}$, $\mathbf{y}_1 \in \mathbb{Z}^{n_1}$, not all zero, such that*

$$\left\| \begin{pmatrix} \mathbf{A}_0\mathbf{x}_0 + \hat{\mathbf{A}}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2 \\ \frac{q}{2}\mathbf{y}_1 - \mathbf{x}_1 \\ -\mathbf{x}_2 \end{pmatrix} + \begin{pmatrix} \mathbf{e} \\ \frac{1}{2}\mathbf{s}_1 \\ \mathbf{s}_2 \end{pmatrix} \mod q \right\| < c \cdot \|\mathbf{e}_{\text{lat}}\|, \tag{5}$$

where each component of the resulting vector is represented by the unique residue in $(-q/2, q/2]$.

If no non-zero vectors satisfy (5), then the geometric separation of Condition (4) holds and the attack succeeds with sufficiently high probability. Consequently, the false positive probability is upper-bounded by the probability that Lemma 8 admits a non-zero solution.

Since the integer vectors range over an infinite set, working directly with (5) is inconvenient. We therefore relax it to a necessary condition: if a false positive occurs, then there must exist a non-zero vector $(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \mathbf{y}_1)$ with components taken from the finite domains $\mathbf{x}_0 \in \mathbb{Z}_q^{n_0}$, $\mathbf{x}_1 \in \mathbb{Z}_q^{n_1}$, $\mathbf{x}_2 \in \mathbb{Z}_q^{n_2}$, $\mathbf{y}_1 \in \mathbb{Z}_2^{n_1}$, such that

$$\left\| \begin{pmatrix} \mathbf{A}_0\mathbf{x}_0 + \hat{\mathbf{A}}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2 \\ \frac{q}{2}\mathbf{y}_1 - \mathbf{x}_1 \\ -\mathbf{x}_2 \end{pmatrix} + \begin{pmatrix} \mathbf{e} \\ \frac{1}{2}\mathbf{s}_1 \\ \mathbf{s}_2 \end{pmatrix} \mod q \right\| < c \cdot \|\mathbf{e}_{\text{lat}}\|. \tag{6}$$

In other words, if no non-zero finite-field vector satisfies (6), then the scores of all incorrect guesses are strictly lower than that of the correct guess, and the attack succeeds. Consequently, estimating the probability of event (6) lies at the heart of the analysis.

To prepare for the subsequent probability estimation, we make the modulo-q operation in (6) explicit. Condition (6) is equivalent to the existence of integer vectors $\mathbf{k}_0 \in \mathbb{Z}^m$, $\mathbf{k}_1 \in \mathbb{Z}^{n_1}$, $\mathbf{k}_2 \in \mathbb{Z}^{n_2}$ such that

$$\begin{cases} \mathbf{A}_0\mathbf{x}_0 + \hat{\mathbf{A}}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2 + \mathbf{e} = q\mathbf{k}_0 + \mathbf{r}_0, \\ \frac{q}{2}\mathbf{y}_1 - \mathbf{x}_1 + \frac{1}{2}\mathbf{s}_1 = q\mathbf{k}_1 + \mathbf{r}_1, \\ -\mathbf{x}_2 + \mathbf{s}_2 = q\mathbf{k}_2 + \mathbf{r}_2, \end{cases} \tag{7}$$

where $\mathbf{r} = (\mathbf{r}_0, \mathbf{r}_1, \mathbf{r}_2)$ satisfies $\|\mathbf{r}\| < c \cdot \|\mathbf{e}_{\text{lat}}\|$, and each component is represented by its unique residue in $(-q/2, q/2]$.

Thus, the probability of a false positive event reduces to the probability that the system (7) admits a non-zero solution. In the following, we estimate this probability by analyzing the expected number of non-zero solutions to (7), splitting the analysis into the general case $(\mathbf{x}_1, \mathbf{x}_2) \neq (\mathbf{0}, \mathbf{0})$ and the degenerate case $(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{0}, \mathbf{0})$. We shall show that the expected number of solutions $(\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \mathbf{y}_1)$ is far smaller than 1, from which it follows that a false positive occurs only with sufficiently low probability.

General Case: $(\mathbf{x_1}, \mathbf{x_2}) \neq (\mathbf{0}, \mathbf{0}) \mod q$ For each arbitrarily fixed triple $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{y}_1)$, define

$$\mathbf{r}_{1,2} = (\mathbf{r}_1, \mathbf{r}_2) = \left(\frac{q}{2}\mathbf{y}_1 - \mathbf{x}_1 + \frac{1}{2}\mathbf{s}_1,\ -\mathbf{x}_2 + \mathbf{s}_2 \right) \mod q,$$

whose Euclidean norm is $\|\mathbf{r}_{1,2}\| = \sqrt{\|\mathbf{r}_1\|^2 + \|\mathbf{r}_2\|^2}$. For brevity, set $R = c \cdot \|\mathbf{e}_{\text{lat}}\|$. According to the false positive condition (7), if this triple is to participate in a false positive event, there must exist some $r < R$ and some $\mathbf{x}_0 \in \mathbb{Z}_q^{n_0}$ such that

$$\|\mathbf{r}_0\| < t,\ \mathbf{r}_0 = (\mathbf{A}_0\mathbf{x}_0 + \hat{\mathbf{A}}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2 + \mathbf{e}) \mod q,\ t = \sqrt{R^2 - r^2}.$$

Note that t is a real number if and only if $r \leq R$; if $\|\mathbf{r}_{1,2}\| \geq R$, this triple cannot satisfy the false positive condition. Hence, in the subsequent summation, we only need to consider those $\mathbf{r}_{1,2}$ with $\|\mathbf{r}_{1,2}\| < R$.

Next we bound the number of triples $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{y}_1)$ that can yield a given $\mathbf{r}_{1,2}$. The following lemma gives a simple but crucial upper bound.

Lemma 9. *For any fixed* $\mathbf{r}_{1,2} = (\mathbf{r}_1, \mathbf{r}_2) \in \frac{1}{2}\mathbb{Z}^{n_1} \times \frac{1}{2}\mathbb{Z}^{n_2}$ *defined by*

$$\mathbf{r}_1 = \frac{q}{2}\mathbf{y}_1 - \mathbf{x}_1 + \frac{1}{2}\mathbf{s}_1 \mod q, \quad \mathbf{r}_2 = -\mathbf{x}_2 + \mathbf{s}_2 \mod q,$$

the number of triples $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{y}_1)$ *with* $\mathbf{x}_1 \in \mathbb{Z}_q^{n_1}$, $\mathbf{x}_2 \in \mathbb{Z}_q^{n_2}$, *and* $\mathbf{y}_1 \in \mathbb{Z}_2^{n_1}$ *that produce this* $\mathbf{r}_{1,2}$ *is at most* 2^{n_1}.

Proof. The condition on $\mathbf{r}_2$ directly determines $\mathbf{x}_2 = \mathbf{s}_2 - \mathbf{r}_2 \mod q$, so $\mathbf{x}_2$ is uniquely fixed. It remains to count the number of pairs $(\mathbf{x}_1, \mathbf{y}_1)$ that satisfy the condition on $\mathbf{r}_1$. For each choice of $\mathbf{y}_1 \in \mathbb{Z}_2^{n_1}$, the equation

$$\mathbf{x}_1 = \frac{q}{2}\mathbf{y}_1 + \frac{1}{2}\mathbf{s}_1 - \mathbf{r}_1 \mod q$$

yields at most one candidate for $\mathbf{x}_1$ in $\mathbb{Z}_q^{n_1}$: if the right-hand side is an integer vector, that vector is the unique candidate; otherwise, no valid $\mathbf{x}_1$ exists for that $\mathbf{y}_1$. Since there are exactly 2^{n_1} possible choices for $\mathbf{y}_1$, the total number of valid triples $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{y}_1)$ is bounded by 2^{n_1}. □

To proceed with the analysis of the norm condition on $\mathbf{r}_0$, we need the joint distribution of the submatrices $\mathbf{A}_0$, $\hat{\mathbf{A}}_1$, and $\mathbf{A}_2$. The following lemma shows that they remain mutually independent.

Lemma 10. *Let q be a prime and let m, n_0, n_1, n_2 be positive integers. Let $\mathbf{A}_0 \in \mathbb{Z}_q^{m\times n_0}$, $\mathbf{A}_1 \in \mathbb{Z}_q^{m\times n_1}$, and $\mathbf{A}_2 \in \mathbb{Z}_q^{m\times n_2}$ be mutually independent random matrices, i.e., all their entries are independent and uniformly distributed over $\mathbb{Z}_q$. Take $\gamma = 2$ and set $\hat{\mathbf{A}}_1 = \gamma\mathbf{A}_1 \mod q$. Then $\mathbf{A}_0$, $\hat{\mathbf{A}}_1$, and $\mathbf{A}_2$ are mutually independent.*

Proof. Since q is an odd prime, multiplication by $\gamma = 2$ is a bijection on $\mathbb{Z}_q$. Hence each entry of $\hat{\mathbf{A}}_1$ remains independent and uniform, and the joint distribution of $(\mathbf{A}_0, \hat{\mathbf{A}}_1, \mathbf{A}_2)$ factorizes. The detailed verification is given in Appendix B. □

Fix $\mathbf{x}_1, \mathbf{x}_2$ and define $\mathbf{u} = \hat{\mathbf{A}}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2 + \mathbf{e} \in \mathbb{Z}_q^m$. When $(\mathbf{x}_1, \mathbf{x}_2) \neq (\mathbf{0}, \mathbf{0})$, Lemma 10 implies that $\hat{\mathbf{A}}_1$ and $\mathbf{A}_2$ are independent, and therefore $\mathbf{u}$ is uniformly distributed over $\mathbb{Z}_q^m$. For a fixed $\mathbf{A}_0$, consider the set spanned by its columns, $L_0 = \{\mathbf{A}_0\mathbf{x}_0 : \mathbf{x}_0 \in \mathbb{Z}_q^{n_0}\} \subset \mathbb{Z}_q^m$. whose size satisfies $|L_0| \leq q^{n_0}$. The probability that there exists $\mathbf{x}_0$ with $\|\mathbf{A}_0\mathbf{x}_0 + \mathbf{u}\| < t$ is bounded by

$$\Pr_{\mathbf{u}}\big(\exists \mathbf{x}_0,\ \|\mathbf{A}_0\mathbf{x}_0 + \mathbf{u}\| < t\big) \leq \frac{|L_0 \cap B(\mathbf{u}, t)|}{q^m} \leq \frac{q^{n_0} M(t)}{q^m},$$

where $M(t)$ denotes the number of integer points in the m-dimensional closed ball of radius t.

Summarizing the above, for any fixed triple $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{y}_1)$, the contribution of the corresponding $\mathbf{r}_{1,2}$ to the false positive probability is at most $2^{n_1} \cdot \frac{q^{n_0}}{q^m} M(t)$, where $t = \sqrt{R^2 - \|\mathbf{r}_{1,2}\|^2}$. To bound $M(t)$, we invoke a standard geometric counting lemma.

Lemma 11. *Let $B(t) = \{\mathbf{x} \in \mathbb{R}^m : \|\mathbf{x}\| \leq t\}$ be the m-dimensional closed ball of radius t, and let $M(t) = |B(t) \cap \mathbb{Z}^m|$. Denote by $V_m = \frac{\pi^{m/2}}{\Gamma(m/2+1)}$ the volume of the m-dimensional unit ball. Then*

$$M(t) \leq V_m \left(t + \frac{\sqrt{m}}{2}\right)^m.$$

Proof. For each integer point $\mathbf{x} \in \mathbb{Z}^m$, construct the unit cube $C_{\mathbf{x}} = \prod_{i=1}^m [x_i - \frac{1}{2}, x_i + \frac{1}{2}]$ centered at $\mathbf{x}$. These cubes have disjoint interiors and volume 1. If $\mathbf{x} \in B(t)$, then for every point $\mathbf{y} \in C_{\mathbf{x}}$ we have $\|\mathbf{y} - \mathbf{x}\| \leq \frac{\sqrt{m}}{2}$, and therefore $\|\mathbf{y}\| \leq t + \frac{\sqrt{m}}{2}$. Hence, each such cube is contained in the enlarged ball $B(t + \frac{\sqrt{m}}{2})$. Summing over all integer points in $B(t)$ gives

$$M(t) = \sum_{\mathbf{x}\in B(t)\cap\mathbb{Z}^m} \mathrm{Vol}(C_{\mathbf{x}}) \leq \mathrm{Vol}\Big(B\Big(t + \frac{\sqrt{m}}{2}\Big)\Big) = V_m \left(t + \frac{\sqrt{m}}{2}\right)^m.$$

□

Combining Lemma 10 and Lemma 11, each fixed $\mathbf{r}_{1,2}$ contributes to the false positive probability at most

$$2^{n_1} \cdot \frac{q^{n_0}}{q^m} \cdot V_m\big(\sqrt{R^2 - \|\mathbf{r}_{1,2}\|^2} + \frac{\sqrt{m}}{2}\big)^m.$$

Summing over all $\mathbf{r}_{1,2}$ with $\|\mathbf{r}_{1,2}\| < R$ yields the preliminary bound

$$p_{\neq} \leq 2^{n_1} q^{n_0-m} V_m \sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ \|\mathbf{r}\|<R}} \left(\sqrt{R^2 - \|\mathbf{r}\|^2} + \frac{\sqrt{m}}{2}\right)^m,$$

where $d = n_1+n_2$. To simplify notation, let $a = \frac{\sqrt{m}}{2}$ and $g(s) = (\sqrt{R^2 - s^2}+a)^m$. Then

$$p_{\neq} \leq 2^{n_1} q^{n_0-m} V_m \sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ \|\mathbf{r}\|<R}} g(\|\mathbf{r}\|). \tag{8}$$

Now we have reduced the problem to estimating a discrete sum over the half-integer lattice $\frac{1}{2}\mathbb{Z}^d$. Because the summand $g(\|\mathbf{r}\|)$ depends on $\mathbf{r}$ only through its norm, the sum does not admit a convenient closed form. We now bound it from above by a continuous integral via a standard volume argument.

We partition the interval $[0, R)$ into a small inner region $[0, r_0)$ and the bulk region $[r_0, R)$, where we take $r_0 = \frac{\sqrt{d}}{2}$. Correspondingly, the sum splits as

$$\sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ \|\mathbf{r}\|<R}} g(\|\mathbf{r}\|) = \sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ \|\mathbf{r}\|<r_0}} g(\|\mathbf{r}\|) + \sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ r_0\leq\|\mathbf{r}\|<R}} g(\|\mathbf{r}\|).$$

For a half-integer point $\mathbf{r}$ with $r_0 \leq \|\mathbf{r}\| < R$, construct the cube $C_{\mathbf{r}} = \prod_{i=1}^d [r_i - 1/4, r_i + 1/4]$ of side length $1/2$ centered at $\mathbf{r}$; its volume is 2^{-d}. For any $\mathbf{x} \in C_{\mathbf{r}}$, the triangle inequality gives $\|\mathbf{r}\| \geq \|\mathbf{x}\| - \frac{\sqrt{d}}{4}$. Since g is non-increasing on $[0, R]$, we have

$$g(\|\mathbf{r}\|) \leq g\left(\|\mathbf{x}\| - \frac{\sqrt{d}}{4}\right).$$

Multiplying both sides by 2^d and integrating over $C_{\mathbf{r}}$ yields

$$g(\|\mathbf{r}\|) \leq 2^d \int_{C_{\mathbf{r}}} g\left(\|\mathbf{x}\| - \frac{\sqrt{d}}{4}\right) d\mathbf{x}.$$

Summing over all $\mathbf{r}$ with $r_0 \leq \|\mathbf{r}\| < R$ and using the disjointness of these cubes as well as the fact that they are all contained in the region $\{\mathbf{x} \in \mathbb{R}^d : r_0 - \frac{\sqrt{d}}{4} \leq \|\mathbf{x}\| < R + \frac{\sqrt{d}}{4}\}$, we obtain

$$\sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ r_0\leq\|\mathbf{r}\|<R}} g(\|\mathbf{r}\|) \leq 2^d \int_{r_0-\frac{\sqrt{d}}{4}\leq\|\mathbf{x}\|<R+\frac{\sqrt{d}}{4}} g\left(\|\mathbf{x}\| - \frac{\sqrt{d}}{4}\right) d\mathbf{x}.$$

Switching to spherical coordinates with $r = \|\mathbf{x}\|$ and letting $u = r - \frac{\sqrt{d}}{4}$, after simplification we arrive at

$$\sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ r_0\le\|\mathbf{r}\|<R}} g(\|\mathbf{r}\|) \le 2^d S_{d-1} \int_0^R g(u) \left(u + \frac{\sqrt{d}}{4}\right)^{d-1} du, \tag{9}$$

where $S_{d-1} = \frac{2\pi^{d/2}}{\Gamma(d/2)}$ is the surface area of the $(d-1)$-dimensional unit sphere. since $r_0 = \frac{\sqrt{d}}{2}$, the lower limit becomes 0.

Let $V_d = \frac{\pi^{d/2}}{\Gamma(d/2+1)}$ denote the volume of the d-dimensional unit ball. For the region $\|\mathbf{r}\| < r_0$, the function values satisfy $g(s) \le g(0)$, To bound the number of half-integer points in this region, we apply the same volume argument as in the bulk part: place a cube of side length $1/2$ around each point; these cubes are disjoint and all lie within a ball of radius $r_0 + \sqrt{d}/4$. Hence, their total volume is at most the volume of that ball, giving

$$\sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ \|\mathbf{r}\|<r_0}} 1 \le 2^d \cdot V_d \cdot \left(r_0 + \frac{\sqrt{d}}{4}\right)^d.$$

Consequently,

$$\sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ \|\mathbf{r}\|<r_0}} g(\|\mathbf{r}\|) \le g(0) \cdot 2^d \cdot V_d \cdot \left(r_0 + \frac{\sqrt{d}}{4}\right)^d. \tag{10}$$

Final Upper Bound for the False Positive Probability Substituting the estimates (9) and (10) into (8), we obtain the following analytic upper bound for the false positive probability in the general case:

$$p_{\neq} \le 2^{n_1} q^{n_0-m} V_m \cdot 2^d \left[S_{d-1} \int_0^R g(u) \left(u + \frac{\sqrt{d}}{4}\right)^{d-1} du \;+\; g(0) V_d \left(\frac{3\sqrt{d}}{4}\right)^d \right],$$

where g, a are as in (8), V_d is the volume of the d-dimensional unit ball, and $S_{d-1} = dV_d$ as defined above.

Special Case: $(\mathbf{x_1}, \mathbf{x_2}) = (\mathbf{0}, \mathbf{0}) \mod q$ It remains to consider the degenerate case where $(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{0}, \mathbf{0}) \mod q$. In this case, the argument that treats $\hat{\mathbf{A}}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2$ as a uniformly random vector no longer applies, and a separate analysis is required. We now show that its contribution, denoted by p_{00}, is likewise sufficiently low under typical parameter choices.

When $\mathbf{x}_1 = \mathbf{x}_2 = \mathbf{0} \mod q$, the system (7) reduces to

$$\frac{q}{2}\mathbf{y}_1 + \frac{1}{2}\mathbf{s}_1 = q\mathbf{k}_1 + \mathbf{r}_1, \qquad \mathbf{s}_2 = q\mathbf{k}_2 + \mathbf{r}_2.$$

From the second equation, since $\|\mathbf{s}_2\| \ll q$ and each component is represented by its unique residue in $(-q/2, q/2]$, we must have $\mathbf{k}_2 = \mathbf{0}$ and $\mathbf{r}_2 = \mathbf{s}_2$. Denote $r_2 = \|\mathbf{s}_2\|$.

For the first equation, rearranging yields

$$\mathbf{r}_1 = \frac{q}{2}\mathbf{y}_1 + \frac{1}{2}\mathbf{s}_1 - q\mathbf{k}_1,$$

where again each component is taken modulo q with the absolutely smallest representative. Since q is odd, $\frac{q}{2}$ is a half-integer. Recall that $\mathbf{y}_1 \in \mathbb{Z}_2^{n_1}$, so each of its coordinates is either 0 or 1. If some coordinate equals 1, then $\frac{q}{2} \cdot 1$ modulo q evaluates to $-q/2$ in the centered representation, whose absolute value is $q/2$. Given that $\|\mathbf{s}_1\|_\infty \ll q$, the corresponding component of $\mathbf{r}_1$ has absolute value at least $q/2 - \frac{1}{2}\|\mathbf{s}_1\|_\infty$. Because the parameters are chosen such that $R = c \cdot \|\mathbf{e}_{\text{lat}}\| \ll q/2$, this value strictly exceeds R, and consequently $\|\mathbf{r}_{1,2}\| \geq \|\mathbf{r}_1\| > R$, which violates Condition (5).

Hence, the only scenario that can possibly satisfy $\|\mathbf{r}_{1,2}\| < R$ is $\mathbf{y}_1 = \mathbf{0}$. In this case, the equation further forces $\mathbf{k}_1 = \mathbf{0}$ and $\mathbf{r}_1 = \frac{1}{2}\mathbf{s}_1$, whose norm we denote by $r_1 = \frac{1}{2}\|\mathbf{s}_1\|$. Thus $\mathbf{r}_{1,2} = (\frac{1}{2}\mathbf{s}_1, \mathbf{s}_2)$ is uniquely determined, and its norm is $r_{12} = \sqrt{r_1^2 + r_2^2}$. The condition $\|\mathbf{r}_{1,2}\| < R$ is equivalent to $r_1^2 + r_2^2 < R^2$, which is clearly satisfied for the typical parameters under consideration.

For this unique combination, the false positive condition additionally requires the existence of a non-zero vector $\mathbf{x}_0 \in \mathbb{Z}_q^{n_0}$ such that

$$\|\mathbf{A}_0\mathbf{x}_0 + \mathbf{e}\| < t, \ t = \sqrt{R^2 - r_{12}^2}.$$

Applying the union bound and the estimate for the number of integer points in a ball from Lemma 11, the probability of this event is at most

$$\Pr(\exists \mathbf{x}_0 \neq \mathbf{0}) \leq (q^{n_0} - 1)\frac{M(t)}{q^m} \leq q^{n_0} \cdot \frac{V_m\left(t + \frac{\sqrt{m}}{2}\right)^m}{q^m}.$$

Therefore, the contribution of the special case to the expected false positive probability satisfies

$$p_{00} \leq q^{n_0 - m} V_m \left(t + \frac{\sqrt{m}}{2}\right)^m.$$

Substituting typical parameter values confirms that p_{00} is far smaller than 1, so the probability of this special case is indeed sufficiently low.

Having bounded both the general and the degenerate case, we now combine them into a single overall bound. The total false positive probability satisfies

$$\Pr[\text{false positive}] \leq p_{\neq} + p_{00}.$$

Under typical parameter choices, both $p_{\neq}$ and p_{00} are sufficiently small, and hence the false-positive probability is upper-bounded by $p_{\neq} + p_{00}$.

Aborted Probability-Decreasing Enumeration of $\mathbf{s_0}$ For the enumeration of $\mathbf{s}_0$, we adopt the aborted probability-decreasing strategy of [5]. Let X be a random variable taking values $x_0, x_1, \ldots, x_{r-1}$ with probabilities $p_k = \Pr(X = x_k)$ in non-increasing order $p_0 \geq p_1 \geq \cdots \geq p_{r-1}$. Then the expected number of guesses to find the correct value is $G(X) = \sum_{i=0}^{r-1} i \cdot p_i$. For an n_0-coordinate secret $\mathbf{s}_0$ with independent coordinates distributed as X, [6] and [5] show that the expected number of guesses in an aborted probability-decreasing strategy has the following closed form

$$G_{\text{class}}(n_0) = \sum_{i=0}^{\mu-1} p'_i \Big(F_i + \frac{f_i(f_i+1)}{2} \Big),$$

where the p'_i are the distinct key probabilities sorted non-increasingly, f_i is the multiplicity of p'_i, and $F_i = \sum_{j=0}^{i-1} f_j$ is the cumulative number of keys with probability strictly larger than p'_i. The number of distinct probabilities is $\mu = \binom{n_0+r-1}{n_0}$, which is polynomial in n_0 for constant r.

If the attacker accepts a failure probability p_{enum} and aborts after the cumulative success probability reaches $1 - p_{\text{enum}}$, the expected number of guesses per round becomes

$$\text{Cost}_{\text{class}}(p_{\text{enum}}) = \sum_{i=0}^{\mu'} p'_i \Big(F_i + \frac{f_i(f_i+1)}{2} \Big) + \Big(1 - \sum_{i=0}^{\mu'} f_i p'_i \Big) F_{\mu'+1},$$

where μ' is the largest index with $\sum_{i=0}^{\mu'} f_i p'_i \leq 1 - p_{\text{enum}}$. If a failure occurs, the process is restarted, giving an expected total cost of at most $1/(1-p_{\text{enum}})$ times the single-round cost. We denote by $\mathbb{G}(n_0, p_{\text{enum}})$ this final expected number of guesses; it replaces the exhaustive factor $|\mathbb{S}_0|$ in the complexity formula.

With the correctness analysis of Section 3 and the probability-decreasing enumeration strategy for $\mathbf{s}_0$ incorporated, we now present the complete procedure of our provable FFT-accelerated dual attack as Algorithm 1.

Extension to larger values of γ. The analysis above is detailed for $\gamma = 2$, but extends naturally to $\gamma = 3$. The key points are that Lemma 10 remains valid when $\gcd(q, \gamma) = 1$ (in particular, for $\gamma = 3$ when q is a power of two), and that the geometric counting grid changes from half-integer to 1/3-spaced points; the subsequent counting and integral-estimation arguments remain applicable with adjusted constants. Details are given in Appendix B.

The above discussion describes the full attack procedure. We now turn to analyzing the time complexity of this procedure. The total cost consists of three components: short vector sampling, table construction for each guess of $\mathbf{s}_0$, and the FFT computation. The complexity of short vector sampling and the FFT computation has already been given above. We now give the formula for the table construction cost for each guess of $\mathbf{s}_0$.

For each guess of $\mathbf{s}_0$ that survives this truncation, the dominant computational step is the construction of a table T that aggregates the contributions

of all N dual vectors, after which an FFT recovers the candidate for $\mathbf{s}_1$. The following lemma, restated in our notation from [18], bounds the cost of building T.

Lemma 12 (Table Construction Cost). *[18] Let $W = (\mathbf{w}_1, \ldots, \mathbf{w}_N)$ be i.i.d. short vectors drawn from the conditional distribution $D^*_{L_q^\perp(\mathbf{B}_{\text{dual}}),qs}$. When enumerating all candidates for $\mathbf{s}_0$, the table T can be built using $T_{\text{table}}(N) = O(N)$ addition operations.*

Algorithm 1: Provable FFT-Accelerated Dual Attack

Input: LWE samples $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$ with $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$; parameters n_0, n_1, n_2 such that $n = n_0 + n_1 + n_2$; distribution parameter $s > 0$; number of short vectors $N \in \mathbb{N}$; enumeration failure probability $p_{\text{enum}} \in (0, 1]$;

Output: Guess $\mathbf{s}_0 \in \mathbb{Z}_q^{n_0}$ for part of the secret.

Split $\mathbf{A}$ into column blocks $\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2$ corresponding to the three parts of the secret.

Choose $\gamma = 2$ if q is prime, or $\gamma = 3$ if q is a power of two. Set $\hat{\mathbf{A}}_1 = \gamma \mathbf{A}_1 \mod q$.

Sample N independent short vectors $\mathbf{v}_j = (\mathbf{w}_j, \hat{\mathbf{v}}_{1,j}, \mathbf{v}_{2,j})$ from $D_{L_q^\perp(\mathbf{B}_{\text{dual}}),qs}$.

Initialize $F_{\max} \leftarrow -\infty$, $\mathbf{s}_0^{\text{best}} \leftarrow \perp$, $\mathbf{s}_1^{\text{best}} \leftarrow \perp$.

for *each candidate $\tilde{\mathbf{s}}_0 \in \mathbb{S}_0$ in probability-decreasing order with abort probability p_{enum} (see [5])* **do**

  Initialize an n_1-dimensional table T of size γ^{n_1} with all entries zero.

  for *each sampled vector $(\mathbf{w}_j, \hat{\mathbf{v}}_{1,j}, \mathbf{v}_{2,j})$* **do**

    $\mathbf{v}_{0,j}^T \leftarrow \mathbf{w}_j^T \mathbf{A}_0$;

    $\hat{\mathbf{v}}_{1,j}^T \leftarrow \mathbf{w}_j^T \hat{\mathbf{A}}_1$;

    $\textit{phase} \leftarrow \exp\left(\frac{2\pi i}{q}\left(\mathbf{w}_j^T \mathbf{b} - \mathbf{v}_{0,j}^T \tilde{\mathbf{s}}_0\right)\right)$;

    $\mathbf{c} \leftarrow \hat{\mathbf{v}}_{1,j} \mod \gamma \in \mathbb{Z}_\gamma^{n_1}$;

    Add *phase* to the cell $T[\mathbf{c}]$.

  Perform n_1-dimensional FFT on T.;

  $F \leftarrow (1/N) \cdot \max_{\tilde{\mathbf{s}}_1 \in \mathbb{Z}_\gamma^{n_1}} \Re\big(T[\tilde{\mathbf{s}}_1]\big)$;

  if $F > F_{\max}$ **then**

    $F_{\max} \leftarrow F$;

    $\left(\mathbf{s}_0^{\text{best}}, \mathbf{s}_1^{\text{best}}\right) \leftarrow (\tilde{\mathbf{s}}_0, \tilde{\mathbf{s}}_1 \mod \gamma)$;

return $\mathbf{s}_0^{\text{best}}$;

Collecting the above cost estimates, we obtain the following theorem that encapsulates both the correctness and complexity of the attack.

Theorem 2 (Correctness and Complexity). *Let $\mathbf{A} = [\mathbf{A}_0 \ \mathbf{A}_1 \ \mathbf{A}_2] \in \mathbb{Z}_q^{m\times n}$ be a uniformly random matrix, let $\mathbf{e}$ be a short noise vector, and let $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \mod q$. Suppose $\mathbf{v}_1, \ldots, \mathbf{v}_N \in L_q^\perp(\mathbf{B}_{\text{dual}})$ are i.i.d. short vectors obtained via*

the sampler of Lemma 3 with parameters satisfying the conditions stated in the analysis. Suppose the attack uses the aborted probability-decreasing enumeration as described in Section 3 with a prescribed failure probability $p_{\text{enum}} \in (0,1)$*. Then Algorithm 1 outputs the correct pair* $(\mathbf{s}_0, \mathbf{s}_1 \mod \gamma)$ *with probability at least*

$$p_{\text{dual}} \geq (1 - p_{\text{enum}}) \left(1 - |\mathbb{S}_0| \cdot \gamma^{n_1} \cdot 2^{-\Omega(N\delta^2)}\right) (1 - p_{\neq} - p_{00}) .$$

The overall time complexity of the algorithm is

$$\begin{aligned} T_{\text{total}} =& T_{\text{SV}}(m + n_1 + n_2, \beta, q^{n_1+n_2}) + O\big(N + \gamma^{n_1}\big) \cdot \mathbb{G}(n_0, p_{\text{enum}}) \\ =& T_{BKZ}(m + n_1 + n_2, \beta) + N \cdot T_{MCMC}\left(L_q^{\perp}(\mathbf{B}_{dual}), qs\right) \\ &+ O\big(N + \gamma^{n_1}\big) \cdot \mathbb{G}(n_0, p_{\text{enum}}), \end{aligned}$$

where $\mathbb{G}(n_0, p_{\text{enum}})$ *denotes the expected number of guesses for* $\mathbf{s}_0$ *under the aborted probability-decreasing enumeration, and* $T_{\text{SV}}(m, \beta, q^n)$ *is the cost of generating a single short vector from Black-Box 1.*

Proof. The correctness follows from Lemmas 6 and 7, together with Corollary 2 and the geometric separation event $\mathcal{G}$, which ensure that the correct target point uniquely achieves the largest score with the stated probability. The probability-decreasing enumeration retains the correct candidate with probability at least $1 - p_{\text{enum}}$, giving the corresponding overall success probability.

For the complexity, generating the N short vectors costs $N \cdot T_{\text{SV}}$, while the lookup-table construction and FFT evaluation over each $\mathbf{s}_0 \in \mathbb{S}_0$ cost $O(N)$ and $O(\gamma^{n_1})$, respectively. Adding basis reduction, sampling, and post-processing costs yields the claimed complexity. □

Remark 1 (Effect of approximate Gaussian sampling). The approximate sampler of Lemma 3 incurs statistical distance at most $N\eta$ over N samples, so the success probability in Theorem 2 degrades by at most $N\eta$. Choosing $\eta \ll 1/N$ makes this degradation negligible.

4 Complexity Estimation

We now instantiate the provable framework of Section 3 with concrete parameters, derive the feasibility constraints, and give numerical cost estimates.

4.1 Parameter Constraints and Feasible Region

Recall the geometric separation constant c defined in Section 3. The attack imposes several conditions on c, which in turn determine the admissible range for the Gaussian width s. Full derivations are deferred to Appendix C; here we state the essential conclusions.

Lower bounds on c. Correctness requires $c > 2$. Additionally, the integral splitting in the false-positive estimate forces $c > \sqrt{n_1 + n_2}/(2\|\mathbf{e}_{\text{lat}}\|)$. Hence

$$c > c_{\text{low}} := \max\Big(2,\ \sqrt{n_1 + n_2}/(2\|\mathbf{e}_{\text{lat}}\|)\Big).$$

For ML-KEM parameters, the second term is $O(1)$, and the dominant condition is $c > 2$.

Upper bounds on c from false-positive control. The false-positive probability upper bound derived in Section 3 can be expressed as a function of the radius $R = c\,\|\mathbf{e}_{\text{lat}}\|$. Specifically, define

$$F(R) \;=\; 2^{n_1} q^{n_0-m} V_m \sum_{\substack{\mathbf{r}\in \frac{1}{2}\mathbb{Z}^d \\ \|\mathbf{r}\|<R}} + q^{n_0-m} V_m \left(t + \frac{\sqrt{m}}{2}\right)^m,$$

where the summand and the parameter t are as specified in the derivation of the false-positive bound. Notice that $F(R)$ is a strictly increasing function of R (all terms grow with R). Two restrictions follow from this fact:

1. For a prescribed negligible threshold κ, define

   $$R_{\max} = \sup\{\, R \geq 0 \mid F(R) \leq \kappa \,\}.$$

   To guarantee Pr[false positive] $\leq \kappa$, we must have $R \leq R_{\max}$, i.e.,

   $$c \leq c_{\max}^{(\text{prob})} := R_{\max}\ /|\mathbf{e}_{\text{lat}}\|.$$

2. The degenerate case $(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{0}, \mathbf{0})$ requires $R < q/2 - \|\mathbf{s}_1\|_\infty/2$, yielding $c < c_{\max}^{(\text{mod})} = (q/2 - \|\mathbf{s}_1\|_\infty/2)/\|\mathbf{e}_{\text{lat}}\|$.

The effective upper bound is therefore

$$c \leq c_{\text{up}} := \min\big(c_{\max}^{(\text{prob})},\ c_{\max}^{(\text{mod})}\big).$$

From c to a lower bound on s. Recall from the correctness analysis that the safety margin satisfies $\tau = \frac{1}{s}\sqrt{\frac{M}{2\pi}},\ M = m + n_1 + n_2$. The separation condition requires c to satisfy $\tau < (c-2)\|\mathbf{e}_{\text{lat}}\|$, which, given a value of c, translates directly into a lower bound on s:

$$s > s_{\min}(c) := \frac{1}{(c-2)\|\mathbf{e}_{\text{lat}}\|}\sqrt{\frac{M}{2\pi}}. \tag{11}$$

This inequality is rigorous and follows solely from the geometric correctness condition; it does not depend on any heuristic.

4.2 Parameter Setup

We now specify how the various parameters are chosen in a concrete LWE instance. Fix n, m, and let q be a prime power. Let $\mathbf{s} \in \mathbb{Z}_q^n$ be the secret and let $\sigma_e > 0$. Let $(\mathbf{A}, \mathbf{b})$ be sampled from $\mathrm{LWE}(m, \mathbf{s}, D_{\mathbb{Z}_q,\sigma_e})$ and write $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$.

For the concrete complexity estimates for ML-KEM, following [21], we use the Gaussian norm estimates

$$\|\mathbf{e}\| \le \sigma_e \sqrt{\frac{m}{2\pi}},\ \|\mathbf{s}_1\| \le \sigma_s \sqrt{\frac{n_1}{2\pi}},\ \|\mathbf{s}_2\| \le \sigma_s \sqrt{\frac{n_2}{2\pi}}, \qquad \sigma_s = \sigma_e = \sqrt{\pi\eta}.$$

Consequently,

$$\|\mathbf{e}_{\mathrm{lat}}\| \le \sqrt{\frac{m\sigma_e^2}{2\pi} + \frac{n_1\sigma_s^2}{8\pi} + \frac{n_2\sigma_s^2}{2\pi}}.$$

Choose a Gaussian width $s > 0$ and set $\delta = \rho_{1/s}(\|\mathbf{e}_{\mathrm{lat}}\|)/10$. We then choose the number of short dual vectors as $N = (\mathrm{poly}(m, n) + n \log q)/\delta^2$, which ensures the required success probability. Recall from Section 4.1 that the correctness condition requires

$$\tau < (c-2)\|\mathbf{e}_{\mathrm{lat}}\|, \qquad \tau = \frac{1}{s}\sqrt{\frac{m + n_1 + n_2}{2\pi}}.$$

In practice, we use a conservative choice of the separation parameter c. For the ML-KEM parameter sets considered in this work, we set $c = 7$.

4.3 Numerical Results

Existing provable dual-attack analyses include [21,22,23], with the best results due to [21]. We use the same complexity estimation as in those works. We employ the same complexity abstractions and cost models as those used in these works for a consistent comparison.

Table 2. Concrete cost estimates for ML-KEM parameters.

Scheme	Cost(bits)	m	n_0	n_1	n_2	$\ln(p_{\neq} + p_{00})$	p_{enum}	β	s	c
ML-KEM-512	216	600	57	84	371	-254.52	0.50	662	0.07	7
ML-KEM-768	317	810	83	122	563	-275.39	0.40	1022	0.09	7
ML-KEM-1024	426	1100	132	157	735	-375.52	0.35	1410	0.09	7

For every entry in Table 2, the selected parameter s satisfies the feasibility constraint (11). Under the common cost estimates, our estimated attack cost improves upon the recomputed estimates of [21] by 22, 30, and 52 bits for ML-KEM-512, ML-KEM-768, and ML-KEM-1024, respectively. These improvements result from the combined effect of the general dual lattice formulation, the FFT-based score computation, and the refined probabilistic treatment developed in this work.

Table 3. Cost estimates for the ML-KEM parameter sets. All costs are logarithms in base two.

Scheme	Attack cost		Reduction
	[21]	Ours	
ML-KEM-512	238	216	22
ML-KEM-768	347	317	30
ML-KEM-1024	478	426	52

5 Comparison with the Contradictory Regime

We compare our provable success regime with the contradictory regime identified in [10]. Unlike [21], which expresses this regime in terms of the Gaussian width s, we formulate it geometrically using the separation parameter c.

The heuristic distinguishability analysis of [18] was revisited in [10], which identified a "contradictory regime" where the dual attack cannot reliably distinguish the correct secret from incorrect guesses. Consequently, the heuristic estimates of [18] and related works [2,8,13] may fail in this regime.

Under the orthogonal-lattice framework of [21], the contradictory regime corresponds to

$$\ln T \le N\varepsilon^2,\ \varepsilon = \exp(-\pi\sigma_e^2\ell^2),\ \ell = \sqrt{4/3}\cdot \mathrm{GH}(L_{\mathrm{dual}}).$$

Using $\|\mathbf{e}_{\mathrm{lat}}\|^2 \approx M\sigma_e^2/(2\pi)$ and $\mathrm{GH}(L_{\mathrm{dual}}) \approx \lambda_1 > c\|\mathbf{e}_{\mathrm{lat}}\|$, this condition implies $c \le T^{1/M}$, and hence heuristically $c < T^{1/M}$. For ML-KEM parameters, $T^{1/M}$ remains close to 1. The detailed derivation and the existence of feasible s satisfying $\ln T \le N\varepsilon^2$ are given in Appendix D.

Thus, when $c \le T^{1/M} \approx 1$, the geometric separation required by our analysis breaks down, corresponding to the contradictory regime of [10]. In contrast, Sections 4.1 show that our attack succeeds with sufficiently high probability for

$$c \in (c_{\mathrm{low}}, c_{\mathrm{up}}],$$

where c_{low} and c_{up} arise from the correctness and false-positive conditions, respectively. For the ML-KEM parameters considered here, c_{low} is roughly 2, leaving the unresolved regime $T^{1/M} \le c \le c_{\mathrm{low}}$.

Table 4. Geometric separation regimes for the parameter c.

Regime	**Range of c**	**Property**
Contradictory	$c < T^{1/M}$	heuristic distinguishability fails
Unresolved	$T^{1/M} \le c \le c_{\mathrm{low}}$	Geometric separation unresolved
Success (this work)	$c_{\mathrm{low}} < c \le c_{\mathrm{up}}$	Correctness and false positives controlled
Not guaranteed	$c > c_{\mathrm{up}}$	Provable false-positive control unavailable

For $c > c_{\mathrm{up}}$, the false-positive bound no longer guarantees a sufficiently low probability, so correctness is not covered by our analysis. For our concrete parameters, e.g., $c = 7$, we have $c_{\mathrm{low}} < c \leq c_{\mathrm{up}}$, placing the attack in the provable success regime.

6 Conclusion

We have presented a provable correctness analysis of the Guo–Johansson-style FFT-based dual attack over the general dual lattice $L_q(\mathbf{A}_{\mathrm{dual}}^T)$. This extends the previous analysis from $L_q^\perp(\mathbf{A}_{\mathrm{dual}})$ and reduces the required number of LWE samples from $m \approx 2n$ to $m \approx n$, while retaining FFT-based score computation.

At the theoretical level, we introduce a geometric separation constant c to characterize the success and failure regimes, and derive an explicit false-positive bound by combining integer-point counting with a modulus-exclusion argument. Together with the high-probability norm bound under the actual CBD distribution of ML-KEM, these results yield a rigorous correctness analysis.

For ML-KEM-512, ML-KEM-768, and ML-KEM-1024, our concrete attack costs are 216, 317, and 426 bits, respectively, improving upon the corresponding recomputed estimates of [21] by 22, 30, and 52 bits under the common probabilistic criterion. These results demonstrate the benefits of combining the general dual-lattice formulation, FFT-based score computation, and refined probabilistic treatment while retaining rigorous correctness guarantees.

References

1. Aharonov, D., Regev, O.: Lattice problems in np /spl cap/ conp. In: 45th Annual IEEE Symposium on Foundations of Computer Science. pp. 362–371 (2004). https://doi.org/10.1109/FOCS.2004.35
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)
4. Bashiri, K., Wiemers, A.: On the independence heuristic in the dual attack. Journal of Mathematical Cryptology **19** (2025)
5. Budroni, A., Mårtensson, E.: Further improvements of the estimation of key enumeration with applications to solving lwe. Cryptography and Communications **16**(5), 1163–1182 (2024). https://doi.org/10.1007/s12095-024-00587-0
6. Budroni, A., Mårtensson, E.: Improved estimation of key enumeration with applications to solving lwe. In: 2023 IEEE International Symposium on Information Theory (ISIT). pp. 495–500 (2023). https://doi.org/10.1109/ISIT54713.2023.10206474
7. Carrier, K., Meyer-Hilfiger, C., Shen, Y., Tillich, J.P.: Assessing the impact of a variant of matzov's dual attack on kyber. In: Advances in Cryptology–CRYPTO 2025. pp. 444–476. Springer Nature Switzerland (2025)
8. Carrier, K., Shen, Y., Tillich, J.P.: Faster dual lattice attacks by using coding theory. IACR Cryptol. ePrint Arch. **2022**, 1750 (2022), https://api.semanticscholar.org/CorpusID:255188792

9. Chevignard, C., Schrottenloher, A., Shen, Y.: Quantum algorithm for discrete gaussian sampling. Cryptology ePrint Archive, Paper 2026/984 (2026), https://eprint.iacr.org/2026/984
10. 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)
11. Ducas, L., Pulles, L.N.: Accurate score prediction for dual-sieve attacks: L. ducas, ln pulles. Journal of Cryptology **39**(1), 8 (2026)
12. Glaser, T., May, A., Nowakowski, J.: Entropy suffices for guessing most keys (2023), https://api.semanticscholar.org/CorpusID:268386912
13. 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)
14. Hastad, J.: Dual vectors and lower bounds for the nearest lattice point problem. pp. 75–81. Combinatorica (1988)
15. Hou, J., Jiang, H.: Careful with the ring: Enhanced hybrid decoding attacks against module/ring-LWE. Cryptology ePrint Archive, Paper 2026/366 (2026), https://eprint.iacr.org/2026/366
16. Laarhoven, T., Walter, M.: Dual lattice attacks for closest vector problems (with preprocessing). In: Paterson, K.G. (ed.) Topics in Cryptology – CT-RSA 2021. pp. 478–502. Springer International Publishing, Cham (2021)
17. Levieil, É., Fouque, P.A.: An improved lpn algorithm. In: International Conference on Security and Cryptography for Networks (SCN'06). pp. 348–359. Springer Berlin Heidelberg (2006)
18. MATZOV: Report on the security of lwe: Improved dual lattice attack. Tech. rep., IACR Cryptol. ePrint Arch. (2022), report 2022/445
19. Micciancio, D., Regev, O.: Lattice-based 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
20. Ogilvie, T.: On the concrete hardness gap between MLWE and LWE. Cryptology ePrint Archive, Paper 2026/279 (2026), https://eprint.iacr.org/2026/279
21. Pouly, A., Shen, Y.: Provable dual attacks on learning with errors. In: Joye, M., Leander, G. (eds.) Advances in Cryptology – EUROCRYPT 2024. pp. 256–285. Springer Nature Switzerland, Cham (2024)
22. Pouly, A., Shen, Y.: Discrete gaussian sampling for bkz-reduced basis. In: Niederhagen, R., Saarinen, M.J.O. (eds.) Post-Quantum Cryptography. pp. 63–88. Springer Nature Switzerland, Cham (2025)
23. 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
24. Qu, H., Xu, G.: On the provable dual attack for lwe by modulus switching. In: Hanaoka, G., Yang, B.Y. (eds.) Advances in Cryptology – ASIACRYPT 2025. pp. 34–64. Springer Nature Singapore, Singapore (2026)
25. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Symposium on the Theory of Computing (2005), https://api.semanticscholar.org/CorpusID:53223958
26. Serre, J.P.: A course in arithmetic (1973), https://api.semanticscholar.org/CorpusID:118597616
27. Wang, R.J., Wang, Z.X., Zheng, Q.X., Zhao, X.: LaMS: A p-adic layered modulus switching for provable dual attacks on LWE. Cryptology ePrint Archive, Paper 2026/1326 (2026), https://eprint.iacr.org/2026/1326

28. 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
29. Zhao, N., Ling, C.: Improved dual attack via quantum rejection sampling. Cryptology ePrint Archive, Paper 2026/979 (2026), https://eprint.iacr.org/2026/979

A Proof of Lemma 4

We provide the detailed proof of the lattice equivalence stated in Lemma 4.

For any $\mathbf{z} \in L_q^{\perp}(\mathbf{B}_{\text{dual}})$, write $\mathbf{z} = \begin{pmatrix} \mathbf{x} \\ \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} \in \mathbb{Z}^{m+n-n_0}$ with $\mathbf{x} \in \mathbb{Z}^m$, $\mathbf{y}_1 \in \mathbb{Z}^{n_1}$, $\mathbf{y}_2 \in \mathbb{Z}^{n_2}$. Then

$$
\begin{aligned}
\mathbf{z} \in L_q^{\perp}(\mathbf{B}_{\text{dual}}) &\iff \mathbf{B}_{\text{dual}}^T \mathbf{z} = \mathbf{0} \mod q \\
&\iff \begin{pmatrix} \hat{\mathbf{A}}_1^T & -\mathbf{I}_{n_1} & \mathbf{0} \\ \mathbf{A}_2^T & \mathbf{0} & -\mathbf{I}_{n_2} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} = \mathbf{0} \mod q \\
&\iff \begin{cases} \hat{\mathbf{A}}_1^T \mathbf{x} - \mathbf{y}_1 = 0 \mod q \\ \mathbf{A}_2^T \mathbf{x} - \mathbf{y}_2 = 0 \mod q \end{cases} \\
&\iff \begin{cases} \hat{\mathbf{A}}_1^T \mathbf{x} = \mathbf{y}_1 \mod q \\ \mathbf{A}_2^T \mathbf{x} = \mathbf{y}_2 \mod q \end{cases} \\
&\iff (\hat{\mathbf{A}}_1 \mid \mathbf{A}_2)^T \mathbf{x} = \begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} \mod q \\
&\iff \mathbf{z} = \begin{pmatrix} \mathbf{x} \\ \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix} \in L_q(\mathbf{A}_{\text{dual}}^T).
\end{aligned}
$$

This completes the proof.

B Proof of Lemma 10

We provide the detailed proof of the independence of $\mathbf{A}_0$, $\hat{\mathbf{A}}_1$, and $\mathbf{A}_2$ stated in Lemma 10.

Since $q > 2$ is prime, $\gamma = 2$ is invertible in the ring $\mathbb{Z}_q$, so the multiplication map $\phi(x) = \gamma x \mod q$ is a bijection. Consequently, each entry of $\hat{\mathbf{A}}_1 = \gamma \mathbf{A}_1 \mod q$ is also an independent uniform random variable over $\mathbb{Z}_q$, and hence $\hat{\mathbf{A}}_1$ is itself a uniform random matrix.

Consider the joint distribution of the triple $(\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2)$. By the independence assumption, all $m(n_0 + n_1 + n_2)$ entries are independent and uniform, so for any matrices $\mathbf{X}_0 \in \mathbb{Z}_q^{m \times n_0}$, $\mathbf{X}_1 \in \mathbb{Z}_q^{m \times n_1}$, $\mathbf{X}_2 \in \mathbb{Z}_q^{m \times n_2}$ we have

$$
\Pr(\mathbf{A}_0 = \mathbf{X}_0,\ \mathbf{A}_1 = \mathbf{X}_1,\ \mathbf{A}_2 = \mathbf{X}_2) = q^{-m(n_0+n_1+n_2)}. \tag{12}
$$

Define the bijection $\Phi(\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2) = (\mathbf{A}_0,\ \gamma\mathbf{A}_1 \bmod \mathrm{q},\ \mathbf{A}_2)$. For arbitrary given matrices $\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2$, using the bijectivity of Φ we obtain

$$\begin{aligned}
&\Pr(\mathbf{A}_0 = \mathbf{B}_0,\ \hat{\mathbf{A}}_1 = \mathbf{B}_1,\ \mathbf{A}_2 = \mathbf{B}_2)\\
=&\Pr(\mathbf{A}_0 = \mathbf{B}_0,\ \gamma\mathbf{A}_1 = \mathbf{B}_1,\ \mathbf{A}_2 = \mathbf{B}_2)\\
=&\Pr(\mathbf{A}_0 = \mathbf{B}_0,\ \mathbf{A}_1 = \gamma^{-1}\mathbf{B}_1,\ \mathbf{A}_2 = \mathbf{B}_2)\\
=&q^{-m(n_0+n_1+n_2)}.
\end{aligned}$$

On the other hand, the marginal distributions are

$$\Pr(\mathbf{A}_0 = \mathbf{B}_0) = q^{-mn_0},\ \Pr(\hat{\mathbf{A}}_1 = \mathbf{B}_1) = q^{-mn_1},$$

and $\Pr(\mathbf{A}_2 = \mathbf{B}_2) = q^{-mn_2}$, whose product also equals $q^{-m(n_0+n_1+n_2)}$. Since the joint probability equals the product of the marginals for all $\mathbf{B}_0, \mathbf{B}_1, \mathbf{B}_2$, the three matrices are mutually independent.

Extension to larger values of γ. The analysis above is presented in detail for $\gamma = 2$, which is the setting used in our concrete instantiation. The underlying arguments, however, are not specific to this choice and extend naturally to other admissible values of γ. In particular, the conclusion of Lemma 10 holds whenever $\gcd(q, \gamma) = 1$. Hence, when q is a power of two, the choice $\gamma = 3$ satisfies this condition. The corresponding independence property therefore remains valid, and the uniformity conclusions used in the subsequent analysis continue to hold without changing the structure of the argument.

The geometric counting argument also extends naturally to $\gamma = 3$. For $\gamma = 2$, the relevant quantities are represented by points on a finite half-integer grid. Replacing $\gamma = 2$ by $\gamma = 3$ changes this grid to one with spacing $1/3$. The resulting set is still finite and discrete, although its size may be exponential in the dimension. Thus, the argument remains a finite lattice-point counting problem, with the underlying grid spacing being the only structural change.

Finally, the integral estimates used to upper-bound the number of relevant points depend on the geometry of the counting region and the spacing of the underlying grid, rather than specifically on $\gamma = 2$. The same integral-estimation procedure can therefore be applied to the $1/3$-spaced grid, with the corresponding scaling factors and constants adjusted accordingly. Consequently, the counting, geometric separation, and probability bounds obtained above extend to $\gamma = 3$. We restrict the detailed derivation to $\gamma = 2$ to keep the main analysis concise.

C Derivation of the Parameter Constraints

In this appendix, we provide detailed derivations of the conditions on the geometric constant $c = \lambda_1(L_q(\mathbf{B}_{\mathrm{dual}}))/\|\mathbf{e}_{\mathrm{lat}}\|$ and the Gaussian width s that are stated in Section 4.1. The notations are those of Sections 3 and 4.

C.1 Lower Bounds on c

Correctness condition. From the geometric analysis in Section 3.1, the distance to the lattice satisfies

$$\mathrm{dist}(\mathbf{e}_{\mathrm{lat}}, L) \le \|\mathbf{e}_{\mathrm{lat}}\|, \qquad \mathrm{dist}(\mathbf{e}_{\mathrm{lat}} + \Delta_{\mathrm{lat}}, L) \ge \lambda_1(L_q(\mathbf{B}_{\mathrm{dual}})) - \|\mathbf{e}_{\mathrm{lat}}\|.$$

For the expected score of the correct guess to dominate that of all incorrect guesses, Lemma 2 and the monotonicity of $\rho_{1/s}$ require (see also the end of Section 3)

$$\|\mathbf{e}_{\mathrm{lat}}\| < \lambda_1(L_q(\mathbf{B}_{\mathrm{dual}})) - \|\mathbf{e}_{\mathrm{lat}}\| - \tau.$$

Rearranging gives $\tau < \lambda_1(L_q(\mathbf{B}_{\mathrm{dual}})) - 2\|\mathbf{e}_{\mathrm{lat}}\|$. With $c = \lambda_1(L_q(\mathbf{B}_{\mathrm{dual}}))/\|\mathbf{e}_{\mathrm{lat}}\|$, this becomes

$$\tau < (c-2)\|\mathbf{e}_{\mathrm{lat}}\|.$$

Since $\tau > 0$, we must have $c > 2$, which is the first lower bound:

$$c > 2. \tag{13}$$

Integral splitting point. In the false-positive estimate (Section 3), the discrete sum over half-integer vectors $\mathbf{r} \in \frac{1}{2}\mathbb{Z}^d$ ($d = n_1 + n_2$) with $\|\mathbf{r}\| < R$ is split at $\delta = \frac{\sqrt{d}}{2}$. The bulk part $[\delta, R)$ is then upper-bounded by a continuous integral. For this bulk part to be non-empty, we need $R > \delta$, i.e.,

$$c\|\mathbf{e}_{\mathrm{lat}}\| > \frac{\sqrt{n_1+n_2}}{2},$$

which yields the second lower bound

$$c > c^{\delta}_{\min} := \frac{\sqrt{n_1+n_2}}{2\|\mathbf{e}_{\mathrm{lat}}\|}. \tag{14}$$

Combining (13) and (14) we obtain

$$c > c_{\mathrm{low}} := \max\left(2,\ \frac{\sqrt{n_1+n_2}}{2\|\mathbf{e}_{\mathrm{lat}}\|}\right). \tag{15}$$

C.2 Upper Bounds on c from False-Positive Control

First upper bound: integral estimate of lattice points. The false-positive event is bounded by the probability that an incorrect guess yields a target point within distance $R = c\|\mathbf{e}_{\mathrm{lat}}\|$ from the lattice. In Section 3, this probability is upper-bounded (up to the degenerate case) by

$$p_{\neq} + p_{00} \le 2^{n_1} q^{n_0-m} V_m \sum_{\substack{\mathbf{r}\in\frac{1}{2}\mathbb{Z}^d \\ \|\mathbf{r}\|<R}} g(\|\mathbf{r}\|) + q^{n_0-m} V_m \left(t + \frac{\sqrt{m}}{2}\right)^m,$$

where $V_m = \frac{\pi^{m/2}}{\Gamma(m/2+1)}$ is the volume of the m-dimensional unit ball, and $g(u) = \left(\sqrt{R^2-u^2}+\frac{\sqrt{m}}{2}\right)^m$. The right-hand side, denoted by $F(R)$, is a strictly increasing function of R (all terms grow with R). For a prescribed negligible threshold ε, define

$$R_{\max} = \sup\{\, R \geq 0 \mid F(R) \leq \varepsilon \,\}.$$

To guarantee $p_{\neq} + p_{00} \leq \varepsilon$, we must have $R \leq R_{\max}$, i.e.,

$$c \leq c_{\max}^{(\mathrm{prob})} := \frac{R_{\max}}{\|\mathbf{e}_{\mathrm{lat}}\|}.$$

In practice, $R_{\max}$ is determined numerically by evaluating $F(R)$ for increasing R; the integral approximation of the sum used in the proof of Lemma 11 provides a smooth upper bound that facilitates this computation.

Second upper bound: modulus exclusion for the degenerate case. The special case $(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{0}, \mathbf{0})$ requires a separate treatment. There, the equations reduce to

$$\frac{q}{2}\mathbf{y}_1 + \frac{1}{2}\mathbf{s}_1 = q\mathbf{k}_1 + \mathbf{r}_1, \qquad \mathbf{s}_2 = q\mathbf{k}_2 + \mathbf{r}_2,$$

with $\mathbf{r}_1, \mathbf{r}_2$ represented by their absolutely smallest residues modulo q. If some component of $\mathbf{y}_1$ is non-zero, the corresponding component of $\frac{q}{2}\mathbf{y}_1$ is congruent to $\pm q/2$ modulo q, and its absolute residue is at least $\frac{q}{2} - \frac{1}{2}\|\mathbf{s}_1\|_\infty$ (after combining with the $\frac{1}{2}\mathbf{s}_1$ term and subtracting the integer multiple $qk_{1,i}$). Thus, for any $\mathbf{y}_1 \neq \mathbf{0}$,

$$\|\mathbf{r}\| \geq \|\mathbf{r}_1\| \geq \frac{q}{2} - \frac{1}{2}\|\mathbf{s}_1\|_\infty.$$

The false-positive condition requires $\|\mathbf{r}\| < R$. Hence, if

$$R < \frac{q}{2} - \frac{1}{2}\|\mathbf{s}_1\|_\infty,$$

then no non-zero $\mathbf{y}_1$ can contribute, and the analysis of this case reduces to the unique configuration $\mathbf{y}_1 = \mathbf{0}$, whose probability is already bounded by the p_{00} term in Section 3. Therefore, to safely ignore non-zero $\mathbf{y}_1$, we impose

$$c < c_{\max}^{(\mathrm{mod})} := \frac{q/2 - \|\mathbf{s}_1\|_\infty/2}{\|\mathbf{e}_{\mathrm{lat}}\|}.$$

Combined upper bound. Both (C.2) and (C.2) must hold, so

$$c \leq c_{\mathrm{up}} := \min\left(c_{\max}^{(\mathrm{prob})},\, c_{\max}^{(\mathrm{mod})}\right). \tag{16}$$

C.3 Feasibility of c and the Lower Bound on s

From (15) and (16), the geometric constant c must satisfy

$$c_{\mathrm{low}} < c \leq c_{\mathrm{up}}.$$

The attack is theoretically feasible whenever this interval is non-empty.

For any such c, the correctness margin condition $\tau < (c-2)\|\mathbf{e}_{\text{lat}}\|$ translates directly into a lower bound on s:

$$\tau = \frac{1}{s}\sqrt{\frac{m+n_1+n_2}{2\pi}} < (c-2)\|\mathbf{e}_{\text{lat}}\| \\ \implies \qquad s > \frac{1}{(c-2)\|\mathbf{e}_{\text{lat}}\|}\sqrt{\frac{m+n_1+n_2}{2\pi}} =: s_{\min}(c). \tag{17}$$

D Detailed Derivations for the Comparison with the Contradictory Regime

This appendix gives a one-sided comparison between the provable geometric regime established in Section 3 and the contradictory condition considered in the prior heuristic framework. The distinguishing lattice in this work is structured, and we do not assume that the Gaussian heuristic predicts its first minimum. Instead, we use Minkowski's first theorem, which gives a rigorous one-sided relation valid for every full-rank lattice.

Throughout this appendix, let

$$L := L_q(B_{\text{dual}}), \qquad M := m+n_1+n_2, \qquad E := \|e_{\text{lat}}\|.$$

Let

$$V_M := \frac{\pi^{M/2}}{\Gamma(M/2+1)}$$

be the volume of the Euclidean unit ball in dimension M, and define the exact volume-heuristic radius by

$$\text{GH}_{\text{vol}}(L) := \left(\frac{\det(L)}{V_M}\right)^{1/M}.$$

This is the quantity denoted by $\text{GH}(L)$ below. If a different asymptotic normalization of the Gaussian heuristic is used, the finite-dimensional conversion factor described in Remark 2 must be included.

D.1 A rigorous one-sided lower bound for the volume-heuristic radius

Let $\mathcal{E}_{\text{sep}}$ be the event that the nonzero-solution condition in inequality (5) has no solution. The false-positive analysis in Section 3 gives

$$\Pr[\mathcal{E}_{\text{sep}}] \geq 1 - p_{\neq} - p_{00}. \tag{18}$$

We also use the parameter restriction from Section 4.1,

$$cE < \frac{q}{2} - \frac{\|s_1\|_\infty}{2}, \tag{19}$$

which in particular implies $cE < q/2$.

Proposition 1. *On the event* $\mathcal{E}_{\mathrm{sep}}$,

$$\lambda_1(L) \geq \sqrt{c^2-1}\, E. \tag{20}$$

Consequently,

$$\Pr\Big[\lambda_1(L) \geq \sqrt{c^2-1}\, E\Big] \geq 1 - p_{\neq} - p_{00}. \tag{21}$$

Proof. Assume for contradiction that

$$\lambda_1(L) < \sqrt{c^2-1}\, E.$$

Let $z \in L \setminus \{0\}$ be a shortest vector. By subtracting a vector in $q\mathbb{Z}^M$ if necessary, z may be chosen in centered form; this operation preserves membership in L and does not increase its norm. Moreover, (19) ensures that the vector considered here is not a nonzero vector of $q\mathbb{Z}^M$, so its residue corresponds to a nonzero choice of the variables in inequality (5), with $x_0 = 0$ and $y_1 = 0$. The same is true for $-z$.

Centered reduction modulo q does not increase the Euclidean norm. Hence

$$\begin{aligned}
\min_{\varepsilon\in\{-1,1\}} \big\|(e_{\mathrm{lat}} + \varepsilon z) \bmod q\big\|^2 &\leq \min_{\varepsilon\in\{-1,1\}} \|e_{\mathrm{lat}} + \varepsilon z\|^2 \\
&= E^2 + \|z\|^2 - 2\big|\langle e_{\mathrm{lat}}, z\rangle\big| \\
&\leq E^2 + \lambda_1(L)^2 \\
&< c^2 E^2.
\end{aligned}$$

Thus one of z and $-z$ gives a nonzero solution to inequality (5), contradicting $\mathcal{E}_{\mathrm{sep}}$. This proves (20); (21) follows from (18).

Minkowski's first theorem gives, for every full-rank M-dimensional lattice,

$$\lambda_1(L) \leq 2\left(\frac{\det(L)}{V_M}\right)^{1/M} = 2\,\mathrm{GH}_{\mathrm{vol}}(L). \tag{22}$$

Combining Proposition 1 with (22) yields the following rigorous one-sided bound:

$$\mathrm{GH}_{\mathrm{vol}}(L) \geq \frac{\sqrt{c^2-1}}{2}\, E \qquad \text{on } \mathcal{E}_{\mathrm{sep}}. \tag{23}$$

No approximation of the form $\lambda_1(L) \approx \mathrm{GH}(L)$ is used in this argument.

Remark 2 (Alternative normalization of the Gaussian heuristic). If the Gaussian heuristic is instead defined by

$$\mathrm{GH}_{\mathrm{asym}}(L) := \sqrt{\frac{M}{2\pi e}}\,\det(L)^{1/M},$$

then

$$\mathrm{GH}_{\mathrm{vol}}(L) = a_M\,\mathrm{GH}_{\mathrm{asym}}(L), \qquad a_M := \frac{V_M^{-1/M}}{\sqrt{M/(2\pi e)}}.$$

Accordingly, (23) becomes

$$\mathrm{GH}_{\mathrm{asym}}(L) \geq \frac{\sqrt{c^2-1}}{2a_M}\, E.$$

The factor a_M tends to 1 as M grows and is close to 1 in the dimensions considered here, but it should be retained in a fully explicit finite-dimensional statement.

D.2 A necessary condition for the contradictory regime

The classical contradictory condition considered in the prior heuristic framework is

$$T^{-1/M}\,\mathrm{GH}_{\mathrm{vol}}(L) \leq \sqrt{\frac{M}{2\pi}}\,\sigma_e, \qquad T := |\mathcal{S}_0|2^{n_1}. \tag{24}$$

Denote the event in (24) by $\mathcal{C}_{\mathrm{contra}}$. The use of this condition is itself part of the prior heuristic framework. In this subsection we derive only a necessary consequence of (24); we do not claim that the converse holds.

Set

$$U := \sqrt{\frac{M}{2\pi}}\,\sigma_e, \qquad \alpha := \frac{E}{U}. \tag{25}$$

On $\mathcal{E}_{\mathrm{sep}}$, combining (24) with (23) gives

$$T^{-1/M}\,\frac{\sqrt{c^2-1}}{2}E \leq U.$$

Equivalently,

$$\mathcal{C}_{\mathrm{contra}} \quad\Longrightarrow\quad c \leq \sqrt{1+\frac{4T^{2/M}}{\alpha^2}} \qquad \text{on } \mathcal{E}_{\mathrm{sep}}. \tag{26}$$

In particular, (26) is a necessary condition, not a sufficient condition and not an equivalence.

For a uniform high-probability formulation, let $\alpha_{\min} > 0$ and define the norm event

$$\mathcal{E}_{\mathrm{norm}}(\alpha_{\min}) := \{E \geq \alpha_{\min} U\}.$$

Suppose an explicit lower-tail estimate for the error and secret distributions gives

$$\Pr[\mathcal{E}_{\mathrm{norm}}(\alpha_{\min})] \geq 1 - \eta_{\mathrm{norm}}. \tag{27}$$

Then, whenever

$$c > c_{\mathrm{contra}} := \sqrt{1+\frac{4T^{2/M}}{\alpha_{\min}^2}}, \tag{28}$$

the events $\mathcal{E}_{\mathrm{sep}}$, $\mathcal{E}_{\mathrm{norm}}(\alpha_{\min})$, and (24) cannot hold simultaneously. Consequently,

$$\Pr[\mathcal{C}_{\mathrm{contra}}] \leq p_{\neq} + p_{00} + \eta_{\mathrm{norm}} \qquad \text{for } c > c_{\mathrm{contra}}. \tag{29}$$

No independence between the separation and norm events is required for this union-bound argument.

If the asymptotic normalization $\mathrm{GH}_{\mathrm{asym}}$ from Remark 2 is used in the classical contradictory condition, the corresponding threshold is

$$c_{\mathrm{contra}}^{(\mathrm{asym})} = \sqrt{1 + \frac{4a_M^2 T^{2/M}}{\alpha_{\min}^2}}.$$

D.3 The condition $\ln T < N\varepsilon^2$ and the Gaussian width

We next derive the dependence of $\ln T < N\varepsilon^2$ on the Gaussian width s, keeping all constants explicit. Let

$$\delta := \frac{1}{10}\rho_{1/s}(E) = \frac{1}{10}\exp(-\pi s^2 E^2).$$

By the explicit Hoeffding and union bounds used in Lemma 7, the probability of a concentration failure over all T candidates is at most

$$2T\exp\left(-\frac{N\delta^2}{2}\right).$$

Thus, for a prescribed concentration failure probability $\eta_{\mathrm{conc}} \in (0,1)$, it is sufficient to choose

$$N = \left\lceil \frac{K}{\delta^2} \right\rceil, \qquad K := 2\ln\left(\frac{2T}{\eta_{\mathrm{conc}}}\right). \tag{30}$$

This replaces an unspecified polynomial factor by an explicit quantity.

Let

$$\varepsilon := \exp(-\pi\sigma_e^2\ell^2), \qquad \ell := \sqrt{\frac{4}{3}}\,\mathrm{GH}_{\mathrm{vol}}(L),$$

as in the prior contradictory-regime formulation. Define

$$N_0 := \frac{K}{\delta^2} \qquad \text{and} \qquad P := \frac{\ln T}{K}.$$

Since $N \geq N_0$, the condition

$$\ln T < N_0\varepsilon^2 \tag{31}$$

is sufficient for $\ln T < N\varepsilon^2$. Moreover,

$$\begin{aligned}
\ln T < N_0\varepsilon^2 &\iff \ln T < \frac{K}{\delta^2}\exp(-2\pi\sigma_e^2\ell^2) \\
&\iff \frac{P}{100}\exp(-2\pi s^2 E^2) < \exp(-2\pi\sigma_e^2\ell^2) \\
&\iff 2\pi s^2 E^2 > 2\pi\sigma_e^2\ell^2 + \ln P - \ln 100.
\end{aligned}$$

Therefore, a sufficient lower bound on s is

$$s > s_{\mathrm{CR}}(E,\ell) := \sqrt{\left[\frac{2\pi\sigma_e^2\ell^2 + \ln P - \ln 100}{2\pi E^2}\right]_+}, \tag{32}$$

where $[x]_+ := \max\{x, 0\}$. In particular, the relevant inequality is a lower bound on s, not an upper bound.

Using the instance-wise ratio $E = \alpha U$ from (25), the same condition can be written as

$$s > \frac{1}{\alpha}\sqrt{\left[\frac{2\pi\ell^2}{M} + \frac{\ln P - \ln 100}{M\sigma_e^2}\right]_+}.$$

Together with the geometric correctness requirement from Section 4.1,

$$s > s_{\min}(c) := \frac{1}{(c-2)E}\sqrt{\frac{M}{2\pi}},$$

a sufficient joint choice is

$$s > \max\{s_{\min}(c), s_{\mathrm{CR}}(E, \ell)\}.$$

Equation (32) shows that, under the sampler abstraction, a sufficiently large finite s satisfies (31). This existence statement does not by itself give a concrete attack cost: by (30),

$$N_0 = 100K \exp(2\pi s^2 E^2),$$

so increasing s can increase the number of required dual vectors exponentially. Concrete parameter claims must therefore substitute the chosen values of s, E, ℓ, and η_{conc} into (30) and include the resulting value of N in the cost estimate.