

A Provable Correctness Analysis of the MATZOV-Style FFT Dual Attack on ML-KEM

Jiale Li¹, Li-Ping Wang¹, and Huaxiong Wang²

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
`hxxwang@ntu.edu.sg`

Abstract. The dual attack is a central tool for evaluating the concrete security of the Learning with Errors (LWE) problem, which underlies lattice-based cryptosystems such as ML-KEM. The MATZOV-style FFT dual attack combines dual distinguishing with an FFT-accelerated guessing phase and provides one of the main frameworks for concrete attack estimates against ML-KEM. However, its correctness analysis has traditionally relied on heuristic independence assumptions, leaving a gap between practical attack estimates and provable guarantees.

In this work, we give a rigorous correctness analysis of a MATZOV-style FFT-based dual attack for the ML-KEM parameter sets. The analysis treats the correct and incorrect candidates separately and applies different conditioning arguments on the two sides, because the former requires a provable lower bound on the score of the true candidate, while the latter requires a uniform upper bound for all wrong recovered guesses. Combining these two parts, we obtain an explicit success criterion and a provable threshold separation for the FFT score.

For concrete cost estimation, we follow the same cost-estimation convention as previous provable dual-attack analyses whenever the same subroutine is involved, using Pouly and Shen’s estimates (EUROCRYPT 2024) as our main baseline. We obtain attack costs of 210, 300, and 410 bits for ML-KEM-512, ML-KEM-768, and ML-KEM-1024, respectively. Relative to this baseline, our estimates reduce the attack costs by 28, 47, and 68 bits for ML-KEM-512, ML-KEM-768, and ML-KEM-1024, respectively.

Keywords: Learning with Errors, Provable dual attack, FFT-based dual attack, ML-KEM.

1 Introduction

The Learning with Errors (LWE) problem, introduced by Regev in 2005 [27], is one of the central foundations of lattice-based cryptography. Its concrete security directly affects the parameter selection of modern post-quantum schemes, including ML-KEM and ML-DSA. Among the known attacks on LWE, the dual

attack has become particularly important in concrete security estimation. Following the introduction of lattice sieving techniques for generating many short dual vectors [3,16] and FFT-based methods for evaluating many guesses simultaneously [17,13], the MATZOV report [20] further optimized the guessing stage and reported highly competitive attack costs against several NIST candidates. These developments made FFT-based dual attacks one of the most influential tools in the concrete analysis of module- and ring-LWE schemes.

The geometric idea behind the dual attack goes back to using short dual vectors to distinguish points close to a lattice from points far away [14]. Aharonov and Regev [1] formalized this approach through decisional Bounded Distance Decoding, and the same principle was later brought into LWE cryptanalysis [21]. In the LWE setting, a short dual vector cancels the secret-dependent part of the samples and leaves a small error term, whereas the corresponding value is close to uniform for random samples. This distributional gap is the basis of the dual distinguisher.

Early dual attacks were too costly for concrete cryptanalysis, mainly because short dual vectors were expensive to obtain. Lattice sieving changed this situation by producing exponentially many short vectors in one run [3], allowing the information from many dual vectors to be aggregated through score functions [16]. Independently, FFT-based evaluation, originally used in coding-theoretic cryptanalysis [17], was adapted to the dual attack by Guo and Johansson [13]. The MATZOV report [20] proposed a different way of incorporating FFT into the dual attack and further optimized the guessing stage. These FFT-based dual attacks substantially reduced the estimated cost for ML-KEM parameters and made the dual attack a central object in concrete security analysis.

Despite their efficiency, these analyses rely on heuristic statistical assumptions. In particular, they often treat score contributions from different short vectors, or scores associated with different secret guesses, as if they were independent. This independence heuristic is delicate. Ducas and Pulles [10,11] showed experimentally that predictions based on such assumptions can deviate significantly from observed attack behavior. Later work further clarified, by explicit covariance computations, that correlations among score terms arise naturally from the shared lattice instance and from dependencies among short vectors [4]. These observations call for a correctness analysis that keeps track of the actual sources of randomness instead of replacing them by an independence model.

A complementary line of work seeks rigorous guarantees for dual attacks without relying on the independence heuristic. Pouly and Shen [24] gave the first provable dual attack for LWE. Qu and Xu later incorporated modulus switching and FFT-based evaluation into this provable framework [26]. Recent preprints have continued to develop this line and related variants, including layered modulus switching, covariance-based score prediction, lossy-source-coding techniques, and dual-enhanced algebraic methods [29,18,19,30].

Other parts of the dual-attack pipeline have been improved. The probability-decreasing key-enumeration strategy used in the MATZOV report has been theoretically justified and refined with early termination [12,6,5]. Different guess-

ing paradigms, including zero-assumption windows and related coding-based approaches, have also been studied [7,8]. In parallel, quantum and structure-aware variants have been proposed for dual attacks and related hybrid strategies [2,32,9,15,22]. These works provide rigorous analyses or refined estimates for important related attacks. However, they do not directly give a correctness proof for the MATZOV-style FFT score function analyzed in this paper.

In this work, we give a provable correctness analysis of a MATZOV-style FFT-based dual attack for the ML-KEM parameter sets. Rather than assuming independence among score terms, we analyze the score function by conditioning on the relevant random sources and treating the correct-candidate and incorrect-candidate sides separately. Our goal is not to re-formalize the full concrete-cost pipeline, but to prove the correctness of the MATZOV-style FFT score under the same LWE abstraction and cost-estimation convention used in prior provable dual-attack analyses. This yields an explicit threshold separation between the true guessed part and all incorrect guessed parts, and provides a rigorous correctness framework for an FFT-based dual attack whose original analysis relied on heuristic independence assumptions.

1.1 Contributions

Our contributions are summarized as follows. We give a correctness analysis for the FFT score used in the MATZOV-style dual attack without assuming independence between score terms. The proof treats the correct-candidate and incorrect-candidate sides separately, reflecting the different lower-bound and upper-bound tasks on the two sides. The resulting bounds lead to an explicit threshold-separation criterion and to concrete cost estimates for ML-KEM parameters.

- **A provable correctness framework for the MATZOV-style FFT dual attack.** We formulate the correctness proof around conditional score averages. After the LWE instance is fixed, concentration only relates the actual score to the conditional mean over the sampled short vectors. One must still control how this conditional mean fluctuates with the remaining instance randomness. Our analysis carries out this second step through moment and covariance estimates, showing that the MATZOV-style separation argument is justified in the parameter regime where the covariance contribution is sufficiently small.
- **Covariance analysis of the FFT score.** We develop a covariance-based analysis of the FFT score for both correct and incorrect candidates. After averaging over the sampled short vector, the relevant quantities become weighted sums over nonzero residue classes modulo q . We organize the cross terms by pairs of residue classes and separate the linearly independent pairs from the linearly dependent ones. This viewpoint lets us combine finite-field Fourier expansion and character orthogonality on the correct-candidate side, while on the incorrect-candidate side it combines the uniformity of the phase over $\mathbb{Z}_q$ with second-moment bounds. The resulting estimates give the covariance and false-positive bounds needed for threshold separation.

- **Concrete cost estimates for ML-KEM.** Using the resulting correctness criterion, and following the same cost-estimation convention as prior provable dual attacks, we obtain costs of 210, 300, and 410 bits for ML-KEM-512, ML-KEM-768, and ML-KEM-1024, respectively. Compared with the corrected ePrint estimates of the previous provable dual-attack baseline, these improve the costs by 28, 47, and 68 bits, respectively.

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

Scheme	Attack cost Reduction		
	[23]	Ours	
ML-KEM-512	238	210	28
ML-KEM-768	347	300	47
ML-KEM-1024	478	410	68

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. For a finite set X , $|X|$ denotes its cardinality, and $\ln$ denotes the natural logarithm. For a complex number z , $\Re(z)$ and $|z|$ denote its real part and modulus, respectively. For $x \in \mathbb{R}$, $\{x\}$ denotes the representative of $x \bmod 1$ in $[-1/2, 1/2)$.

We write $\Pr[\cdot]$ for probability and $\mathbb{E}_x[\cdot]$, $\text{Var}_x(\cdot)$, and $\text{Cov}_x(\cdot, \cdot)$ for expectation, variance, and covariance with respect to the random source x . We use nested conditional expectations, variances, and covariances in the standard sense: the conditioning symbol “ $|\cdot$ ” applies only to the probability operator in which it appears, and outer probability operators act on their own listed random sources, treating conditioned sources as fixed.

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 $\text{LWE}_{n,m,q,\chi_s,\chi_e}$ 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 $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q)$.

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 $\text{LWE}_{n,m,q,\chi_s,\chi_e}$.
- Search-LWE: Given a sample from $\text{LWE}_{n,m,q,\chi_s,\chi_e}$, recover $\mathbf{s}$.

2.2 Discrete Fourier Transform

The Fourier transform is an operation that, for a function $f : G \rightarrow \mathbb{C}$ 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 $f : (\mathbb{Z}/q\mathbb{Z})^n \rightarrow \mathbb{C}$ 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

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

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:

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

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

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 $f_{L,s} : \mathbb{R}^n \rightarrow \mathbb{R}$ 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 [28], Chapter VII, Proposition 15). *For an n -dimensional lattice L , let $s > 0$ and $\mathbf{t} \in \mathbb{R}^n$. Then the following holds:*

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

By the Poisson summation formula, the normalized periodic Gaussian function can be written as the characteristic function of the discrete Gaussian distribution over the dual lattice:

$$f_{L,s}(\mathbf{t}) := \frac{\rho_s(L + \mathbf{t})}{\rho_s(L)} = \mathbb{E}_{\mathbf{w} \leftarrow D_{\tilde{L}, 1/s}} \left[e^{2\pi i \langle \mathbf{t}, \mathbf{w} \rangle} \right].$$

The following inequality relates the value of the periodic Gaussian to the distance from the point to the lattice, and will be used to lower bound the correct-candidate score.

Lemma 2 ([24]). *Let $L \subset \mathbb{R}^m$ be a lattice and let $s > 0$. Then, for every $x \in \mathbb{R}^m$, $f_{L, 1/s}(x) \geq \rho_{1/s}(\text{dist}(x, L))$.*

2.5 Probability Inequalities

Theorem 1 (Hoeffding's inequality). *Let $Z_1, \dots, Z_N$ be independent random variables such that $a_i \leq Z_i \leq b_i$ for each i . Let $S_N = Z_1 + \dots + Z_N$. Then, for every $t > 0$,*

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

Theorem 2 (One-sided Cantelli inequality). *Let Z be a real random variable with $\mathbb{E}[Z] = \mu$, $\text{Var}(Z) = V < \infty$. Then, for every $t > 0$,*

$$\Pr(Z \leq \mu - t) \leq \frac{V}{V + t^2}.$$

Theorem 3 (Markov's inequality). *Let Z be a non-negative random variable with $\mathbb{E}[Z] < \infty$. Then, for every $t > 0$,*

$$\Pr(Z \geq t) \leq \frac{\mathbb{E}[Z]}{t}.$$

Equivalently, for every random variable Y and every $k > 0$,

$$\Pr(|Y| \geq t) \leq \frac{\mathbb{E}[|Y|^k]}{t^k}.$$

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 [31]). *There is an algorithm that, given a basis $\mathbf{B}$ of a lattice $L \subset \mathbb{R}^m$, a parameter $s > 0$, and an accuracy $\xi_{\text{samp}} > 0$, outputs a distribution $D_{L,s,\xi_{\text{samp}}}$ satisfying*

$$d_{\text{TV}}(D_{L,s,\xi_{\text{samp}}}, D_{L,s}) \leq \xi_{\text{samp}}.$$

Its running time is

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

where $\bar{\mathbf{b}}_1, \dots, \bar{\mathbf{b}}_m$ are the Gram-Schmidt vectors of $\mathbf{B}$.

We generate the dual short vectors as follows. Given a block size $2 \leq \beta \leq m$, we compute a basis of $L_q^\perp(\mathbf{A})$, reduce it by BKZ with block size β , and then generate N independent samples using the MCMC Gaussian sampler of Lemma 3 with parameter qs . The sampling cost is written as

$$T_{\text{sampling}}(N) = T_{\text{BKZ}}(m, \beta) + N \cdot T_{\text{MCMC}}(L_q^\perp(\mathbf{A}), qs),$$

where

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

Here $\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_m$ are the Gram-Schmidt vectors of the BKZ- β reduced basis.

3 Provable Correctness Analysis

This section describes the MATZOV-style FFT dual attack considered in this paper and proves its correctness under the stated parameter conditions.

3.1 Formal Description of the MATZOV Algorithm

The secret is divided into three parts: the directly enumerated part s_{guess} with n_{guess} coordinates, the FFT-enumerated part s_{fft} with n_{fft} coordinates, and the part used to construct the dual lattice, s_{dual} , with n_{dual} coordinates. The algorithm first enumerates s_{guess} . For each candidate value, it uses the FFT to compute the scores of all possible s_{fft} , and then determines correctness according to a threshold.

Partition the secret $\mathbf{s}$ and the matrix $\mathbf{A}$ by columns into three parts:

$$\mathbf{s}^T = (\mathbf{s}_{\text{guess}}^T, \mathbf{s}_{\text{fft}}^T, \mathbf{s}_{\text{dual}}^T), \quad \mathbf{A} = (\mathbf{A}_{\text{guess}} \mathbf{A}_{\text{fft}} \mathbf{A}_{\text{dual}}).$$

Define the dual lattice

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

The algorithm proceeds as follows:

Algorithm 1 FFT-Based Dual-Lattice Attack Algorithm

Require: LWE samples $(\mathbf{A}, \mathbf{b})$, parameters $(\beta, N, n_{\text{guess}}, n_{\text{fft}}, n_{\text{dual}}, p, C)$

Ensure: The secret part s_{guess} ; the remaining parts can be recovered recursively

- 1: Construct the dual lattice $\Lambda = L_q^\perp(\mathbf{A}_{\text{dual}}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \mathbf{A}_{\text{dual}} \equiv \mathbf{0} \pmod{q}\}$.
 - 2: Run the short-vector sampling algorithm; see Section 2.6. Obtain N accepted short vectors $\mathbf{w}_j$ satisfying $\mathbf{w}_j \pmod{q} \neq \mathbf{0}$.
 - 3: Compute $\mathbf{v}_{j,\text{guess}} = \mathbf{w}_j^T \mathbf{A}_{\text{guess}}$, $\mathbf{v}_{j,\text{fft}} = \mathbf{w}_j^T \mathbf{A}_{\text{fft}}$, $\mathbf{w}_j^T \mathbf{b}$.
 - 4: **for** each possible $\tilde{\mathbf{s}}_{\text{guess}}$, in decreasing order of probability **do**
 - 5: Initialize an n_{fft} -dimensional array T of size $p^{n_{\text{fft}}}$ with all entries set to zero.
 - 6: **for** $j = 1$ to N **do**
 - 7: Compute $\Delta = \frac{p}{q} (\mathbf{w}_j^T \mathbf{b} - \mathbf{v}_{j,\text{guess}}^T \tilde{\mathbf{s}}_{\text{guess}})$.
 - 8: Determine the index $\text{idx} = \left\lfloor \frac{p}{q} \mathbf{v}_{j,\text{fft}} \right\rfloor \pmod{p}$.
 - 9: Add $e^{2\pi i \Delta / p}$ to $T[\text{idx}]$.
 - 10: **end for**
 - 11: Apply an n_{fft} -dimensional FFT modulo p to T , obtaining $\hat{T}[\tilde{\mathbf{s}}_{\text{fft}}]$.
 - 12: **for** each $\tilde{\mathbf{s}}_{\text{fft}} \in (\mathbb{Z}/p\mathbb{Z})^{n_{\text{fft}}}$ **do**
 - 13: Compute the score $F(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) = \Re(\hat{T}[\tilde{\mathbf{s}}_{\text{fft}}])$.
 - 14: **if** $F(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) > C$ **then**
 - 15: **Output** $\tilde{\mathbf{s}}_{\text{guess}}$ as the correct guess.
 - 16: **end if**
 - 17: **end for**
 - 18: **end for**
-

3.2 Correctness Analysis of the Algorithm

Our analysis follows a conditioning-based approach. Short-vector sampling provides many independent samples, whereas the LWE instance itself is given only

once. Therefore, the score terms are not necessarily unconditionally independent. We first fix the complete LWE instance $(\mathbf{A}, \mathbf{s}, \mathbf{e})$ and apply concentration inequalities only to the sampling randomness of $(\mathbf{w}_1, \dots, \mathbf{w}_N)$. We then study the resulting conditional mean as the remaining random sources vary, and use several good events to separate the correct candidate from all incorrect candidates.

Throughout this section, once a random source no longer appears in the expression under consideration, we omit it from the conditioning and from the subscript of the corresponding probability operator. Thus each expectation, variance, or covariance is taken only over the random sources that still appear in the displayed expression, unless stated otherwise.

Classification of Correct and Incorrect Guesses. Before analyzing the score function, we clarify what is meant by a correct or incorrect candidate. The attack enumerates candidate pairs

$$(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}),$$

but only $\tilde{\mathbf{s}}_{\text{guess}}$ is recovered and used in the recursive step. Hence a candidate is considered incorrect only when

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

while $\tilde{\mathbf{s}}_{\text{fft}}$ is an auxiliary variable. In particular, if

$$\tilde{\mathbf{s}}_{\text{guess}} = \mathbf{s}_{\text{guess}}, \quad \tilde{\mathbf{s}}_{\text{fft}} \neq \mathbf{s}_{\text{fft}},$$

then the candidate may have a different auxiliary part, but it still does not lead to a wrong recovered value.

Therefore, the proof has two parts. On the correct side, it suffices to show that the true pair $(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})$ has score above the threshold. On the wrong side, we show that for every $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$ and every auxiliary choice $\tilde{\mathbf{s}}_{\text{fft}}$, the corresponding score stays below the threshold with high probability.

Distribution of dual short vectors. The attacker first samples short vectors from the dual lattice $\Lambda = L_q^\perp(\mathbf{A}_{\text{dual}})$ with Gaussian parameter $\sigma = qs$, using the sampling procedure described in Section 2.6.

Vectors satisfying $\mathbf{w} \bmod q = \mathbf{0}$ do not contribute to the dual distinguisher and are discarded. Throughout the analysis, N denotes the number of accepted short vectors $\mathbf{w}_1, \dots, \mathbf{w}_N$, namely those satisfying $\mathbf{w}_j \bmod q \neq \mathbf{0}$.

Equivalently, the sampled vectors are distributed according to $D_{\Lambda, \sigma}$ conditioned on $\mathbf{w} \bmod q \neq \mathbf{0}$. Under exact sampling, obtaining N accepted vectors requires $N/(1 - p_0)$ raw MCMC samples in expectation.

All expectations over $\mathbf{w}_j$ are taken with respect to this conditional sampling distribution determined by the fixed value of $\mathbf{A}_{\text{dual}}$.

The residue class $\mathbf{w}_j \bmod q$ lies in $V = \{\mathbf{v} \in \mathbb{Z}_q^m : \mathbf{A}_{\text{dual}}^T \mathbf{v} = \mathbf{0}\}$. Moreover, $|V| = q^{m - \text{rank}(\mathbf{A}_{\text{dual}})}$. Since $\mathbf{A}_{\text{dual}}$ has full column rank with probability

$$p_{\text{rank}} := \prod_{i=0}^{n_{\text{dual}}-1} (1 - q^{i-m}),$$

we have $|V| = q^{m-n_{\text{dual}}}$ on this event. For each $\mathbf{v} \in V$, define

$$p_{\mathbf{v}} := \Pr_{\mathbf{w} \leftarrow D_{\Lambda, \sigma}} [\mathbf{w} \bmod q = \mathbf{v} \mid \mathbf{A}_{\text{dual}}].$$

After conditioning on $\mathbf{w} \bmod q \neq \mathbf{0}$, the corresponding normalized weight is

$$p_{\mathbf{v}}^* := \Pr_{\mathbf{w} \leftarrow D_{\Lambda, \sigma}} [\mathbf{w} \bmod q = \mathbf{v} \mid \mathbf{A}_{\text{dual}}, \mathbf{w} \bmod q \neq \mathbf{0}] = \frac{p_{\mathbf{v}}}{1 - p_{\mathbf{0}}}, \quad \mathbf{v} \in V^{\times} := V \setminus \{\mathbf{0}\}.$$

Score Function and Conditional Expectation. For every candidate

$$(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) \in \{-\eta, \dots, \eta\}^{n_{\text{guess}} + n_{\text{fft}}},$$

define the individual score

$$X_j(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) = \cos\left(\frac{2\pi}{q} \left(\mathbf{w}_j^T \mathbf{b} - \mathbf{w}_j^T \mathbf{A}_{\text{guess}} \tilde{\mathbf{s}}_{\text{guess}} - \frac{q}{p} \left[\frac{p}{q} \mathbf{w}_j^T \mathbf{A}_{\text{fft}} \right] \tilde{\mathbf{s}}_{\text{fft}} \right)\right),$$

where $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$, and p is the modulus-switching parameter. The FFT score used to rank the candidate is therefore

$$F_{\mathbf{W}}(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) = \sum_{j=1}^N \cos\left(\frac{2\pi}{q} \left(\mathbf{w}_j^T \mathbf{b} - \mathbf{w}_j^T \mathbf{A}_{\text{guess}} \tilde{\mathbf{s}}_{\text{guess}} - \frac{q}{p} \left[\frac{p}{q} \mathbf{w}_j^T \mathbf{A}_{\text{fft}} \right] \tilde{\mathbf{s}}_{\text{fft}} \right)\right).$$

When $(\mathbf{A}, \mathbf{s}, \mathbf{e})$ is fixed, the only remaining randomness in X_j comes from the sampled short vector $\mathbf{w}_j$. Hence, by conditioning on the residue class $\mathbf{w}_j \bmod q$, we obtain

$$\mathbb{E}_{\mathbf{w}_j} [X_j(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) \mid (\mathbf{A}, \mathbf{s}, \mathbf{e})] = \sum_{\mathbf{v} \in V^{\times}} p_{\mathbf{v}}^* X_{\mathbf{v}}(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}), \quad j \in [N],$$

where $X_{\mathbf{v}}$ denotes the same score term after fixing $\mathbf{w}_j \bmod q = \mathbf{v}$.

To keep the formulas readable, we often suppress the candidate from the notation. Thus, in the correct-candidate analysis, X_j and $X_{\mathbf{v}}$ stand for

$$X_j(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}), \quad X_{\mathbf{v}}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}),$$

whereas in the incorrect-candidate analysis they stand for

$$X_j(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}), \quad X_{\mathbf{v}}(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}).$$

Correct-Candidate Analysis We first analyze the conditional expectation

$$\mathbb{E}_{\mathbf{w}_j} [X_j \mid (\mathbf{A}, \mathbf{s}, \mathbf{e})]$$

for the correct candidate. Fix a sampled vector $\mathbf{w}_j$. Using $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$, we have

$$\frac{p}{q} \mathbf{w}_j^T \mathbf{b} = \frac{p}{q} \mathbf{w}_j^T (\mathbf{A}_{\text{guess}} \mathbf{s}_{\text{guess}} + \mathbf{A}_{\text{fft}} \mathbf{s}_{\text{fft}} + \mathbf{A}_{\text{dual}} \mathbf{s}_{\text{dual}} + \mathbf{e}) \pmod{p}. \quad (1)$$

Since $\mathbf{w}_j \in L_q^\perp(\mathbf{A}_{\text{dual}})$, the term $\frac{p}{q}\mathbf{w}_j^T \mathbf{A}_{\text{dual}} \mathbf{s}_{\text{dual}}$ is zero modulo p . Moreover,

$$\frac{p}{q}\mathbf{w}_j^T \mathbf{A}_{\text{fft}} \mathbf{s}_{\text{fft}} = \left\lfloor \frac{p}{q}\mathbf{w}_j^T \mathbf{A}_{\text{fft}} \right\rfloor \mathbf{s}_{\text{fft}} + \left\{ \frac{p}{q}\mathbf{w}_j^T \mathbf{A}_{\text{fft}} \right\} \mathbf{s}_{\text{fft}}.$$

Substituting this identity into (1), the phase of the correct-candidate score becomes

$$\frac{p}{q}\mathbf{w}_j^T \mathbf{b} - \frac{p}{q}\mathbf{w}_j^T \mathbf{A}_{\text{guess}} \mathbf{s}_{\text{guess}} - \left\lfloor \frac{p}{q}\mathbf{w}_j^T \mathbf{A}_{\text{fft}} \right\rfloor \mathbf{s}_{\text{fft}} = \left\{ \frac{p}{q}\mathbf{w}_j^T \mathbf{A}_{\text{fft}} \right\} \mathbf{s}_{\text{fft}} + \frac{p}{q}\mathbf{w}_j^T \mathbf{e} \pmod{p}.$$

The following independence lemma will be used repeatedly in the moment calculations.

Lemma 4. *Let q be a prime. Let $\mathbf{A}_{\text{guess}} \in \mathbb{Z}_q^{m \times n_{\text{guess}}}$, $\mathbf{A}_{\text{fft}} \in \mathbb{Z}_q^{m \times n_{\text{fft}}}$ be independent and uniformly distributed. Let $\mathbf{w} \in \mathbb{Z}_q^m \setminus \{\mathbf{0}\}$ be any random vector independent of $\mathbf{A}_{\text{guess}}$ and $\mathbf{A}_{\text{fft}}$. Then $\mathbf{w}$, $\mathbf{w}^T \mathbf{A}_{\text{guess}}$, and $\mathbf{w}^T \mathbf{A}_{\text{fft}}$ are mutually independent. In particular, for every fixed nonzero $\mathbf{w}$,*

$$(\mathbf{w}^T \mathbf{A}_{\text{guess}}, \mathbf{w}^T \mathbf{A}_{\text{fft}}) \sim U(\mathbb{Z}_q^{n_{\text{guess}}} \times \mathbb{Z}_q^{n_{\text{fft}}}),$$

and this distribution does not depend on the value of $\mathbf{w}$.

Proof. A full proof is given in Appendix A.

Correct-Candidate Mean Fix the complete instance $(\mathbf{A}, \mathbf{s}, \mathbf{e})$. Then the short vectors $\mathbf{w}_1, \dots, \mathbf{w}_N$ are independent and identically distributed, and each individual score satisfies $X_j \in [-1, 1]$. Hence the conditional Hoeffding inequality controls the deviation of the actual score $F\mathbf{w} = \sum_{j=1}^N X_j$ from its conditional mean $\sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j \mid \mathbf{A}, \mathbf{s}, \mathbf{e}]$.

This conditional mean, however, still depends on $\mathbf{A}_{\text{fft}}$ and $\mathbf{s}_{\text{fft}}$, and is not directly computable. We therefore first average it over $\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}$ and the short-vector sampling, while keeping $(\mathbf{A}_{\text{dual}}, \mathbf{e})$ fixed. This gives the reference value that will later be compared with the fixed-instance conditional mean via Cantelli's inequality.

For the correct candidate, using Lemma 4, the expectation decomposes as

$$\begin{aligned} & \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}} [F\mathbf{w}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) \mid (\mathbf{A}_{\text{dual}}, \mathbf{e})] \\ &= N \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_j} \left[\exp \left(\frac{2\pi i}{q} \left(\frac{q}{p} \left\{ \frac{p}{q}\mathbf{w}_j^T \mathbf{A}_{\text{fft}} \right\} \mathbf{s}_{\text{fft}} + \mathbf{w}_j^T \mathbf{e} \right) \right) \mid (\mathbf{A}_{\text{dual}}, \mathbf{e}) \right] \\ &= N \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} \left[\exp \left(\frac{2\pi i}{p} \left\{ \frac{p}{q}\mathbf{w}_j^T \mathbf{A}_{\text{fft}} \right\} \mathbf{s}_{\text{fft}} \right) \right] \cdot \mathbb{E}_{\mathbf{w}_j} \left[\exp \left(\frac{2\pi i}{q} \mathbf{w}_j^T \mathbf{e} \right) \mid (\mathbf{A}_{\text{dual}}, \mathbf{e}) \right]. \end{aligned} \tag{2}$$

Let $q' = \frac{q}{\gcd(p, q)}$. For $z \in \mathbb{Z}_q$, the representative of $\left\{ \frac{p}{q}z \right\}$ is taken from the set

$$\mathcal{I}_{q'} := \left\{ \frac{-q' + 1}{2q'} + c_{q'}, \frac{-q' + 3}{2q'} + c_{q'}, \dots, \frac{q' - 1}{2q'} + c_{q'} \right\},$$

where $c_{q'} = 0$ if q' is odd, and $c_{q'} = 1/(2q')$ otherwise. For $s_t \sim \chi_s$, define

$$g_{s_t}(z) := \exp\left(\frac{2\pi i}{p} \left\{\frac{p}{q}z\right\} s_t\right).$$

For $k \in \mathbb{Z}_q$, let

$$c_k(s_t) := \frac{1}{q} \sum_{x \in \mathbb{Z}_q} g_{s_t}(x) \omega_q^{-kx}.$$

Then the Fourier expansion of g_{s_t} over $\mathbb{Z}_q$ is

$$g_{s_t}(z) = \sum_{k \in \mathbb{Z}_q} c_k(s_t) \omega_q^{kz}.$$

In particular, the zero-frequency coefficient is

$$c_0(s_t) = \frac{1}{q} \sum_{x \in \mathbb{Z}_q} \exp\left(\frac{2\pi i}{p} \left\{\frac{p}{q}x\right\} s_t\right).$$

Since $\left\{\frac{p}{q}x\right\}$ is uniformly distributed over $\mathcal{I}_{q'}$ as x ranges over $\mathbb{Z}_q$, and each value appears $\gcd(p, q)$ times, we have

$$\begin{aligned} c_0(s_t) &= \frac{1}{q'} \sum_{\ell=0}^{q'-1} \exp\left[\frac{2\pi i}{p} \left(\frac{-q' + 1 + 2\ell}{2q'} + c_{q'}\right) s_t\right] \\ &= \exp\left(\frac{2\pi i c_{q'}}{p} s_t\right) \frac{\sin(\pi s_t/p)}{q' \sin(\pi s_t/(pq'))}. \end{aligned}$$

When $s_t = 0$, the above expression is interpreted by continuity, and $c_0(0) = 1$.

Now let $\alpha_{\mathbf{v}} := \prod_{t=1}^{n_{\text{fft}}} g_{s_t}(\mathbf{v}^T \mathbf{a}_t)$, where $\mathbf{a}_t$ denotes the t -th column of $\mathbf{A}_{\text{fft}}$. We first give the following important lemma:

Lemma 5 (Character orthogonality over $\mathbb{Z}_q$). *If $\mathbf{a} \sim \mathcal{U}(\mathbb{Z}_q^m)$, then for any $\mathbf{u} \in \mathbb{Z}_q^m$,*

$$\mathbb{E}_{\mathbf{a}} \left[\omega_q^{\mathbf{u}^T \mathbf{a}} \right] = \begin{cases} 1, & \mathbf{u} = \mathbf{0}, \\ 0, & \mathbf{u} \neq \mathbf{0}. \end{cases} \quad (3)$$

Proof. Write $\mathbf{u} = (u_1, \dots, u_m)$, $\mathbf{a} = (a_1, \dots, a_m)$. By the independence between coordinates, we have

$$\mathbb{E}_{\mathbf{a}} \left[\omega_q^{\mathbf{u}^T \mathbf{a}} \right] = \prod_{\ell=1}^d \left(\frac{1}{q} \sum_{x=0}^{q-1} \omega_q^{u_{\ell} x} \right).$$

When $u_{\ell} = 0$, the ℓ -th factor equals 1; when $u_{\ell} \neq 0$, the finite geometric series gives

$$\sum_{x=0}^{q-1} \omega_q^{u_{\ell} x} = \frac{1 - \omega_q^{u_{\ell} q}}{1 - \omega_q^{u_{\ell}}} = 0.$$

Therefore, the product is nonzero if and only if all $u_{\ell} = 0$.

When $\mathbf{v} \neq \mathbf{0}$, by the Fourier expansion and Lemma 5, we obtain $\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}}[\alpha_{\mathbf{v}}] = (\mathbb{E}_{s_t}[c_0(s_t)])^{n_{\text{fft}}}$. Therefore define

$$D_{\text{fft}} := \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} \left[\exp \left(\frac{2\pi i}{p} \left\{ \frac{p}{q} \mathbf{w}_j^T \mathbf{A}_{\text{fft}} \right\} \mathbf{s}_{\text{fft}} \right) \right] = (\mathbb{E}_{s_t}[c_0(s_t)])^{n_{\text{fft}}}.$$

Substituting $c_0(s_t)$ above gives the computable formula

$$D_{\text{fft}} = \left(\mathbb{E}_{s_t} \left[\exp \left(\frac{2\pi i c_{q'}}{p} s_t \right) \frac{\sin(\pi s_t/p)}{q' \sin(\pi s_t/(pq'))} \right] \right)^{n_{\text{fft}}}.$$

Similarly, define

$$\kappa_{\text{sq}} := \left(\mathbb{E}_{s_t} \left[\exp \left(\frac{4\pi i c_{q'}}{p} s_t \right) \left(\frac{\sin(\pi s_t/p)}{q' \sin(\pi s_t/(pq'))} \right)^2 \right] \right)^{n_{\text{fft}}},$$

and

$$\kappa_{\text{abs}} := \left(\mathbb{E}_{s_t} \left[\left(\frac{\sin(\pi s_t/p)}{q' \sin(\pi s_t/(pq'))} \right)^2 \right] \right)^{n_{\text{fft}}}.$$

If χ_s is symmetric under $s_t \mapsto -s_t$, then the above expectations are real and can be written as

$$D_{\text{fft}} = \left(\mathbb{E}_{s_t} \left[R(s_t) \cos \left(\frac{2\pi c_{q'} s_t}{p} \right) \right] \right)^{n_{\text{fft}}},$$

$$\kappa_{\text{sq}} = \left(\mathbb{E}_{s_t} \left[R(s_t)^2 \cos \left(\frac{4\pi c_{q'} s_t}{p} \right) \right] \right)^{n_{\text{fft}}}, \quad \kappa_{\text{abs}} = (\mathbb{E}_{s_t}[R(s_t)^2])^{n_{\text{fft}}},$$

where $R(0) = 1$, and for $s_t \neq 0$,

$$R(s_t) = \frac{\sin(\pi s_t/p)}{q' \sin(\pi s_t/(pq'))}.$$

For $\mathbf{v} \in V$, define $\theta_{\mathbf{v}} := \frac{2\pi}{q} \mathbf{v}^T \mathbf{e}$. Thus $e^{i\theta_{\mathbf{v}}}$ is the noise-related complex exponential factor, whereas $\alpha_{\mathbf{v}}$ contains the contribution from the FFT rounding term. Define

$$\mathcal{I} := \{(\mathbf{v}, \mathbf{w}) \in (V^\times)^2 : \mathbf{v}, \mathbf{w} \text{ are linearly independent over } \mathbb{Z}_q\}, \quad \mathcal{D} := (V^\times)^2 \setminus \mathcal{I}.$$

Write $p_{\text{max}}^* := \max_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^*$ and $W_{\mathcal{D}} := \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{D}} p_{\mathbf{v}}^* p_{\mathbf{w}}^*$. Then the following lemma holds.

Lemma 6. $W_{\mathcal{D}} \leq (q-1)p_{\text{max}}^*$.

Proof. If $\mathbf{v}, \mathbf{w} \in V^\times$ are linearly dependent, then there exists $\lambda \in \mathbb{Z}_q^*$ such that $\mathbf{w} = \lambda \mathbf{v}$. Therefore

$$W_{\mathcal{D}} = \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{I}} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \leq \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* \sum_{\lambda \in \mathbb{Z}_q^*} p_{\lambda \mathbf{v}}^* \leq (q-1)p_{\text{max}}^* \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* = (q-1)p_{\text{max}}^*.$$

To compute $\mathbb{E}_{\mathbf{w}_j} \left[e^{\frac{2\pi i}{q} \mathbf{w}_j^T \mathbf{e}} \middle| (\mathbf{A}_{\text{dual}}, \mathbf{e}) \right]$, we use Lemma 1. For $p_{\mathbf{v}}^*$ defined above, define

$$h^*(\mathbf{e}) := \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* e^{i\theta_{\mathbf{v}}}.$$

Let $L := q\hat{\Lambda}$. Then the following lemma holds.

Lemma 7. *Let $\mathbf{w}_j \leftarrow D_{\Lambda, q_S}$ conditioned on $\mathbf{w}_j \bmod q \neq \mathbf{0}$. Then*

$$\mathbb{E}_{\mathbf{w}_j} \left[e^{\frac{2\pi i}{q} \mathbf{w}_j^T \mathbf{e}} \middle| (\mathbf{A}_{\text{dual}}, \mathbf{e}) \right] = h^*(\mathbf{e}) = \frac{f_{L,1/s}(\mathbf{e}) - p_{\mathbf{0}}}{1 - p_{\mathbf{0}}} \in \mathbb{R}.$$

Proof. The proof follows from the Poisson summation formula and the conditioning $\mathbf{w}_j \bmod q \neq \mathbf{0}$; see Appendix B.

From Lemma 7, we have $h^*(\mathbf{e}) = \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* \cos \theta_{\mathbf{v}}$. Thus the following identities hold.

Lemma 8. $\sum_{\mathbf{v}, \mathbf{w} \in V^\times} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} + \theta_{\mathbf{w}}) = \sum_{\mathbf{v}, \mathbf{w} \in V^\times} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} - \theta_{\mathbf{w}}) = (h^*(\mathbf{e}))^2$.

Proof. This follows by expanding $h^*(\mathbf{e})^2$ and $h^*(\mathbf{e})\overline{h^*(\mathbf{e})}$; see Appendix C.

For $\mathbf{v}, \mathbf{w} \in V^\times$, define

$$H^+(\mathbf{v}, \mathbf{w}) := \mathbb{E}_{\mathbf{a}, s_t} [g_{s_t}(\mathbf{v}^\top \mathbf{a}) g_{s_t}(\mathbf{w}^\top \mathbf{a})], \quad H^-(\mathbf{v}, \mathbf{w}) := \mathbb{E}_{\mathbf{a}, s_t} [g_{s_t}(\mathbf{v}^\top \mathbf{a}) \overline{g_{s_t}(\mathbf{w}^\top \mathbf{a})}].$$

Lemma 9 (Single-column products). *For any $\mathbf{v}, \mathbf{w} \in V^\times$, we have*

$$H^+(\mathbf{v}, \mathbf{w}) = \sum_{\substack{0 \leq k, \ell < q \\ k\mathbf{v} + \ell\mathbf{w} = \mathbf{0}}} \mathbb{E}_{s_t} [c_k(s_t) c_\ell(s_t)], \quad H^-(\mathbf{v}, \mathbf{w}) = \sum_{\substack{0 \leq k, \ell < q \\ k\mathbf{v} - \ell\mathbf{w} = \mathbf{0}}} \mathbb{E}_{s_t} [c_k(s_t) \overline{c_\ell(s_t)}].$$

Furthermore, if $\mathbf{v}, \mathbf{w}$ are linearly independent over $\mathbb{Z}_q$, then

$$H^+(\mathbf{v}, \mathbf{w}) = \mathbb{E}_{s_t} [c_0(s_t)^2], \quad H^-(\mathbf{v}, \mathbf{w}) = \mathbb{E}_{s_t} [|c_0(s_t)|^2].$$

Proof. By the Fourier expansion of g_{s_t} , we obtain

$$\begin{aligned} g_{s_t}(\mathbf{v}^\top \mathbf{a}) g_{s_t}(\mathbf{w}^\top \mathbf{a}) &= \sum_{k, \ell=0}^{q-1} c_k(s_t) c_\ell(s_t) \omega_q^{(k\mathbf{v} + \ell\mathbf{w})^\top \mathbf{a}}, \\ g_{s_t}(\mathbf{v}^\top \mathbf{a}) \overline{g_{s_t}(\mathbf{w}^\top \mathbf{a})} &= \sum_{k, \ell=0}^{q-1} c_k(s_t) \overline{c_\ell(s_t)} \omega_q^{(k\mathbf{v} - \ell\mathbf{w})^\top \mathbf{a}}. \end{aligned}$$

Since finite sums can be exchanged with expectations, Lemma 5 implies that the expectation over $\mathbf{a}$ keeps only the terms satisfying $k\mathbf{v} + \ell\mathbf{w} = \mathbf{0}$ in the first sum and $k\mathbf{v} - \ell\mathbf{w} = \mathbf{0}$ in the second sum. This gives the two displayed formulas. If $\mathbf{v}, \mathbf{w}$ are linearly independent, both equations have only the solution $k = \ell = 0$, proving the last claim.

Lemma 10. For any $\mathbf{v}, \mathbf{w} \in V^\times$, we have

$$|H^+(\mathbf{v}, \mathbf{w})| \leq 1, \quad |H^-(\mathbf{v}, \mathbf{w})| \leq 1.$$

Moreover,

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \alpha_{\mathbf{w}}] = H^+(\mathbf{v}, \mathbf{w})^{n_{\text{fft}}}, \quad \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \overline{\alpha_{\mathbf{w}}}] = H^-(\mathbf{v}, \mathbf{w})^{n_{\text{fft}}}.$$

If $(\mathbf{v}, \mathbf{w}) \in \mathcal{I}$, then

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \alpha_{\mathbf{w}}] = \kappa_{\text{sq}}, \quad \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \overline{\alpha_{\mathbf{w}}}] = \kappa_{\text{abs}}.$$

Proof. The proof is deferred to Appendix D.

Define $\mu(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) := \mathbb{E}_{\mathbf{w}_j} [X_j(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) \mid (\mathbf{A}, \mathbf{s}, \mathbf{e})]$. From the above calculations, we obtain the following expression for the conditional mean under the correct candidate.

Lemma 11. Fix $\mathbf{A}_{\text{dual}}$ and $\mathbf{e}$. For the true candidate $(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})$, we have

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\mu(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) \mid (\mathbf{A}_{\text{dual}}, \mathbf{e})] = D_{\text{fft}} \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* \cos \theta_{\mathbf{v}}.$$

Equivalently,

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) \mid (\mathbf{A}_{\text{dual}}, \mathbf{e})] = N D_{\text{fft}} h^*(\mathbf{e}).$$

Proof. By the definition of D_{fft} , for any $\mathbf{v} \in V^\times$, we have $\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}}] = D_{\text{fft}}$. Therefore

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\mu(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})] = D_{\text{fft}} \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* \cos \theta_{\mathbf{v}}.$$

Then, using

$$h^*(\mathbf{e}) = \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* \cos \theta_{\mathbf{v}}$$

and $F_{\mathbf{W}} = \sum_{j=1}^N X_j$, we obtain the second equality.

First Moment and Cross Second Moment of the Correct-Guess Score

In the previous subsection, after fixing $(\mathbf{A}_{\text{dual}}, \mathbf{e})$, we gave the expression for the expectation of the correct-candidate total score $F_{\mathbf{W}} = \sum_{j=1}^N X_j$ with respect to the random sources $\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_1, \dots, \mathbf{w}_N$. This subsection further studies whether, after fixing the concrete values of $\mathbf{A}_{\text{fft}}$ and $\mathbf{s}_{\text{fft}}$, the conditional mean under short-vector sampling,

$$\sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j \mid \mathbf{A}, \mathbf{s}, \mathbf{e}],$$

is still close to this computable average. To this end, we compute the first moment and the cross second moment of this conditional mean with respect to $\mathbf{A}_{\text{fft}}$ and

$\mathbf{s}_{\text{fft}}$, and use them to control the bad event where it deviates too much from its average.

For $\mathbf{v} \in V^\times$, define

$$Y_{\mathbf{v}} := \Re(e^{i\theta_{\mathbf{v}}} \alpha_{\mathbf{v}}).$$

Thus, after fixing the complete instance $(\mathbf{A}, \mathbf{s}, \mathbf{e})$, the conditional mean of the individual score X_j under short-vector sampling is

$$\mu(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) := \mathbb{E}_{\mathbf{w}_j} [X_j | (\mathbf{A}, \mathbf{s}, \mathbf{e})] = \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* Y_{\mathbf{v}}. \quad (4)$$

Since $F_{\mathbf{W}} = \sum_{j=1}^N X_j$, its variance with respect to $\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}$, conditioned on $(\mathbf{A}_{\text{dual}}, \mathbf{e})$, can be decomposed as

$$\begin{aligned} \text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} (F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})) &= \sum_{j,k=1}^N \text{Cov}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_j, \mathbf{w}_k} (X_j, X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})) \\ &= \sum_{j,k=1}^N \left(\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_j, \mathbf{w}_k} [X_j X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \right. \\ &\quad \left. - \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_j} [X_j | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_k} [X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \right). \end{aligned} \quad (5)$$

On the one hand,

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_j} [X_j | (\mathbf{A}_{\text{dual}}, \mathbf{e})] = \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\mu(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) | (\mathbf{A}_{\text{dual}}, \mathbf{e})],$$

whose computable formula has already been given by Lemma 11. On the other hand, for $j \neq k$, using (4) and the independence of $\mathbf{w}_j$ and $\mathbf{w}_k$ conditioned on $(\mathbf{A}_{\text{dual}}, \mathbf{e})$, we obtain

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_j, \mathbf{w}_k} [X_j X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})] = \sum_{\mathbf{v}, \mathbf{w} \in V^\times} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [Y_{\mathbf{v}} Y_{\mathbf{w}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})]. \quad (6)$$

Therefore, the key remaining task is to evaluate or bound

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [Y_{\mathbf{v}} Y_{\mathbf{w}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})]$$

for pairs $(\mathbf{v}, \mathbf{w}) \in (V^\times)^2$. For linearly independent pairs, Theorem 4 gives an explicit expression; the linearly dependent pairs are controlled through their total weight in Theorem 5.

Theorem 4 (Cross Second Moment of Linearly Independent Pairs). *If $(\mathbf{v}, \mathbf{w}) \in \mathcal{I}$, then*

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [Y_{\mathbf{v}} Y_{\mathbf{w}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] = \frac{\kappa_{\text{sq}}}{2} \cos(\theta_{\mathbf{v}} + \theta_{\mathbf{w}}) + \frac{\kappa_{\text{abs}}}{2} \cos(\theta_{\mathbf{v}} - \theta_{\mathbf{w}}). \quad (7)$$

Proof. Using $\Re(z_1)\Re(z_2) = \frac{1}{2}\Re(z_1 z_2) + \frac{1}{2}\Re(z_1 \bar{z}_2)$ with $z_1 = e^{i\theta_{\mathbf{v}}} \alpha_{\mathbf{v}}$ and $z_2 = e^{i\theta_{\mathbf{w}}} \alpha_{\mathbf{w}}$, we get

$$\begin{aligned} & \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [Y_{\mathbf{v}} Y_{\mathbf{w}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \\ &= \frac{1}{2} \Re \left(e^{i(\theta_{\mathbf{v}} + \theta_{\mathbf{w}})} \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \alpha_{\mathbf{w}}] \right) + \frac{1}{2} \Re \left(e^{i(\theta_{\mathbf{v}} - \theta_{\mathbf{w}})} \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \bar{\alpha}_{\mathbf{w}}] \right). \end{aligned}$$

Since Lemma 10 gives $\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \alpha_{\mathbf{w}}] = \kappa_{\text{sq}}$ and $\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \bar{\alpha}_{\mathbf{w}}] = \kappa_{\text{abs}}$, the claimed identity follows.

Define

$$\begin{aligned} Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) &:= \frac{\kappa_{\text{sq}}}{2} \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{I}} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} + \theta_{\mathbf{w}}) \\ &\quad + \frac{\kappa_{\text{abs}}}{2} \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{I}} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} - \theta_{\mathbf{w}}). \end{aligned} \quad (8)$$

Also define

$$\varepsilon(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) := \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{D}} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [Y_{\mathbf{v}} Y_{\mathbf{w}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})].$$

The following theorem decomposes the cross second moment for $j \neq k$.

Theorem 5. *For any $j \neq k$,*

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_j, \mathbf{w}_k} [X_j X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})] = Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) + \varepsilon(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}).$$

Moreover,

$$|\varepsilon(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})| \leq (q-1)p_{\text{max}}^*.$$

Proof. Split the sum in (6) into the two parts indexed by $\mathcal{I}$ and $\mathcal{D}$. Applying Theorem 4 to the part over $\mathcal{I}$ gives $Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})$, while the remaining part is exactly $\varepsilon(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})$. Since $|Y_{\mathbf{v}}| \leq 1$, we have $|\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [Y_{\mathbf{v}} Y_{\mathbf{w}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})]| \leq 1$. Therefore

$$|\varepsilon(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})| \leq \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{D}} p_{\mathbf{v}}^* p_{\mathbf{w}}^* = W_{\mathcal{D}}.$$

The result follows from Lemma 6.

For any $j \neq k$, Theorem 5 gives the cross second moment, while Lemma 11 gives the corresponding first moment. Substituting these two formulas into the covariance decomposition yields the following bound.

Theorem 6. *Let $\bar{\mu}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) = \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\mu(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) | (\mathbf{A}_{\text{dual}}, \mathbf{e})]$. When $j \neq k$, we have*

$$\begin{aligned} & \text{Cov}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_j, \mathbf{w}_k} (X_j, X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})) \\ &= Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) - (\bar{\mu}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}))^2 + \varepsilon(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}). \end{aligned} \quad (9)$$

Furthermore,

$$\left| \text{Cov}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{w}_j, \mathbf{w}_k} (X_j, X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})) - \left(\frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} - D_{\text{fft}}^2 \right) (h^*(\mathbf{e}))^2 \right| \leq 2(q-1)p_{\text{max}}^*.$$

Proof. The covariance decomposition, Lemma 11, and Theorem 5 first give

$$\text{Cov} = Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) - (\bar{\mu}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}))^2 + \varepsilon(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}).$$

Next, Lemma 8 shows that replacing the sums over $\mathcal{I}$ by the corresponding sums over all pairs in $(V^\times)^2$ introduces an error at most $W_{\mathcal{D}}$. The remaining error term ε is also bounded by $W_{\mathcal{D}}$ by Theorem 5. Therefore,

$$\left| \text{Cov} - \left(\frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} - D_{\text{fft}}^2 \right) (h^*(\mathbf{e}))^2 \right| \leq 2W_{\mathcal{D}} \leq 2(q-1)p_{\max}^*,$$

where the last inequality follows from Lemma 6. The full expansion is given in Appendix E.

Next, we give an upper bound for $p_{\max}^*$. Recall that on the full-rank event of $\mathbf{A}_{\text{dual}}$, which occurs with probability $p_{\text{rank}} = \prod_{i=0}^{n_{\text{dual}}-1} (1 - q^{i-m})$, we have $|V| = q^{m-n_{\text{dual}}}$ and $\det(\Lambda) = q^{n_{\text{dual}}}$. The latter identity will be used in the Poisson-summation estimate below. Define

$$B := \left(\frac{1 + e^{-a}}{1 - e^{-a}} \right)^m, \quad a = \pi s^2.$$

Lemma 12 (Upper Bound on Coset Probability). *Let q be prime, and let $\mathbf{A}_{\text{dual}}$ have full column rank. Then, for any $\mathbf{v} \in V$, $p_{\mathbf{v}} \leq B/|V|$. If furthermore $B < |V|$, then for any $\mathbf{v} \in V^\times$, $p_{\mathbf{v}}^* \leq \frac{B}{|V|-B}$. Therefore*

$$p_{\max}^* = \max_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* \leq \frac{B}{|V|-B} =: \bar{p}_{\max}^*.$$

Proof. The proof is deferred to Appendix F.

$$\text{Define } \Gamma := \left| \left(\frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} - D_{\text{fft}}^2 \right) (h^*(\mathbf{e}))^2 \right| + 2(q-1)\bar{p}_{\max}^*.$$

Theorem 7. *The variance $\text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}}(F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e}))$ has the upper bound $N + N(N-1)\Gamma$.*

Proof. The variance is decomposed into diagonal and off-diagonal terms. The diagonal terms are bounded by N since $X_j \in [-1, 1]$, while the $N(N-1)$ off-diagonal covariance terms are each bounded by Γ by Theorem 6. The full expansion is given in Appendix G.

As explained earlier, after fixing the complete instance, the conditional Hoeffding inequality only shows that the actual score $F_{\mathbf{W}}$ is close to the conditional mean under short-vector sampling, $\sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j | (\mathbf{A}, \mathbf{s}, \mathbf{e})]$. The moment bounds above allow us to compare this fixed-instance conditional mean with its computable average $\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})]$.

We now complete the correct-candidate analysis in two steps: first we control the bad event where the conditional mean is too far from this average; then,

outside this bad event, we apply Theorem 1 to show that the actual score remains close to the fixed-instance conditional mean.

Step 1: Conditional-mean deviation. After fixing $(\mathbf{A}_{\text{dual}}, \mathbf{e})$, the quantity

$$\sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j | (\mathbf{A}, \mathbf{s}, \mathbf{e})]$$

is a random variable with respect to $\mathbf{A}_{\text{fft}}$ and $\mathbf{s}_{\text{fft}}$. By the law of total variance and Theorem 7, its variance is upper bounded by

$$\begin{aligned} & \text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} \left[\sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j | (\mathbf{A}, \mathbf{s}, \mathbf{e})] \middle| (\mathbf{A}_{\text{dual}}, \mathbf{e}) \right] \\ & \leq \text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} (F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})) \leq N + N(N-1)\Gamma. \end{aligned} \quad (10)$$

Taking

$$t = \frac{1}{4} \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})]$$

in Theorem 2, we obtain

$$\begin{aligned} & \Pr \left[\sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j | (\mathbf{A}, \mathbf{s}, \mathbf{e})] \leq \frac{3}{4} \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \right] \\ & \leq \frac{N + N(N-1)\Gamma}{N + N(N-1)\Gamma + \frac{1}{16} (\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})])^2} \\ & = \frac{16 + 16(N-1)\Gamma}{16 + 16(N-1)\Gamma + ND_{\text{fft}}^2 (h^*(\mathbf{e}))^2}. \end{aligned}$$

Denote this bad event by

$$\mathcal{B}_1 := \left\{ \sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j | (\mathbf{A}, \mathbf{s}, \mathbf{e})] \leq \frac{3}{4} \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \right\}.$$

Step 2: Sampling deviation. After fixing $(\mathbf{A}, \mathbf{s}, \mathbf{e})$, the variables $X_1, \dots, X_N$ are independent and satisfy $|X_j| \leq 1$. Hence Theorem 1 gives

$$\begin{aligned} & \Pr \left[F_{\mathbf{W}} - \sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j | (\mathbf{A}, \mathbf{s}, \mathbf{e})] \leq -\frac{1}{4} \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \right] \\ & \leq \exp \left(-\frac{(\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})])^2}{32N} \right) = \exp \left(-\frac{ND_{\text{fft}}^2 (h^*(\mathbf{e}))^2}{32} \right). \end{aligned}$$

Denote this bad event by

$$\mathcal{B}_2 := \left\{ F_{\mathbf{W}} - \sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j | (\mathbf{A}, \mathbf{s}, \mathbf{e})] \leq -\frac{1}{4} \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \right\}.$$

Combining the two steps. By the union bound, $\mathcal{B}_1^c \cap \mathcal{B}_2^c$ occurs with probability at least

$$p_{\text{cor}} = 1 - \frac{16 + 16(N-1)\Gamma}{16 + 16(N-1)\Gamma + ND_{\text{fft}}^2 (h^*(\mathbf{e}))^2} - \exp\left(-\frac{ND_{\text{fft}}^2 (h^*(\mathbf{e}))^2}{32}\right).$$

On this event, the correct-candidate score satisfies

$$\begin{aligned} F_{\mathbf{W}} &> \sum_{j=1}^N \mathbb{E}_{\mathbf{w}_j} [X_j | (\mathbf{A}, \mathbf{s}, \mathbf{e})] - \frac{1}{4} \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \\ &\geq \frac{3}{4} \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] - \frac{1}{4} \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \\ &= \frac{1}{2} \cdot \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} [F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \\ &= \frac{1}{2} \cdot ND_{\text{fft}} h^*(\mathbf{e}). \end{aligned}$$

Thus, with probability at least p_{cor} , $F_{\mathbf{W}} > \frac{1}{2}ND_{\text{fft}}h^*(\mathbf{e})$. It remains to replace $h^*(\mathbf{e})$ by a computable lower bound.

Concrete lower bound on $h^*(\mathbf{e})$. From the definition of $h^*(\mathbf{e})$ in Lemma 7 and Lemma 2, we obtain

$$h^*(\mathbf{e}) = \frac{f_{L,1/s}(\mathbf{e}) - p_0}{1 - p_0} \geq \frac{\rho_{1/s}(\|\mathbf{e}\|) - p_0}{1 - p_0}.$$

Since $\rho_{1/s}(r)$ is decreasing in $r \geq 0$, an upper estimate on $\|\mathbf{e}\|$ gives a lower bound for $h^*(\mathbf{e})$. For the concrete complexity estimates for ML-KEM, where $\mathbf{e} \sim CBD_{\eta}^m$, we follow [24] and use

$$\|\mathbf{e}\| \approx \sigma_e \sqrt{\frac{m}{2\pi}}, \quad \sigma_e = \sqrt{\pi\eta}.$$

Thus we take $u_* := \sqrt{m\eta/2}$.

The value p_0 is also not directly computable from the parameters alone. By Lemma 12, we have $p_0 \leq p_{\text{max}} := B/|V|$. Since the function $p \mapsto (\rho_{1/s}(u_*) - p)/(1 - p)$ is decreasing for $p < \rho_{1/s}(u_*)$, we get, on the norm good event,

$$h^*(\mathbf{e}) \geq \frac{\rho_{1/s}(u_*) - p_0}{1 - p_0} \geq \frac{\rho_{1/s}(u_*) - p_{\text{max}}}{1 - p_{\text{max}}} =: h^*(u_*).$$

Consequently, the correct-candidate score lower bound can be written using only computable parameter quantities.

Lemma 13 (Monotonicity of the Correct-Side Success Probability).
 Let $x := (h^*(\mathbf{e}))^2$ and denote

$$\lambda := \left| \frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} - D_{\text{fft}}^2 \right|, \quad \rho := 2(q-1)p_{\text{max}}^*.$$

Then $\Gamma(x) := \lambda x + \rho$ and

$$p_{\text{cor}}(x) := 1 - \frac{16 + 16(N-1)\Gamma(x)}{16 + 16(N-1)\Gamma(x) + ND_{\text{fft}}^2 x} - \exp\left(-\frac{ND_{\text{fft}}^2 x}{32}\right).$$

Then $p_{\text{cor}}(x)$ is monotonically increasing for $x \geq 0$. Therefore, on the norm good event $\|\mathbf{e}\| \leq u_*$, since $h^*(\mathbf{e}) \geq h^*(u_*)$, we obtain

$$p_{\text{cor}}\left((h^*(\mathbf{e}))^2\right) \geq p_{\text{cor}}\left((h^*(u_*))^2\right).$$

Proof. The proof is given in Appendix H.

The preceding argument completes the correct-candidate side: outside the two bad events $\mathcal{B}_1$ and $\mathcal{B}_2$, the true candidate has score at least $\frac{1}{2}ND_{\text{fft}}h^*(u_*)$, and the probability of these bad events is explicitly bounded by p_{cor} .

Next, we turn to incorrect candidates. The goal is now the opposite one: for every candidate pair satisfying $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$, the score should stay below the same threshold in the success region.

3.3 Analysis of Incorrect Candidates

In the incorrect-candidate analysis, we first fix $\mathbf{A}_{\text{dual}}$ and the true value of $\mathbf{s}_{\text{guess}}$. We then consider only candidates satisfying $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$, so that

$$\delta = \mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{0}.$$

All bounds in this subsection are proved pointwise for such candidates, and the union bound is taken only over these incorrect candidates.

Fix a nonzero residue class $\mathbf{v} \in V^\times$ of the short vector modulo q . Substituting

$$\mathbf{b} = \mathbf{A}_{\text{guess}}\mathbf{s}_{\text{guess}} + \mathbf{A}_{\text{fft}}\mathbf{s}_{\text{fft}} + \mathbf{A}_{\text{dual}}\mathbf{s}_{\text{dual}} + \mathbf{e} \pmod{q}$$

into the phase of the individual score, and denoting $\delta := \mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$, the part related to $\mathbf{A}_{\text{fft}}$ can be written as

$$\mathbf{v}^T \mathbf{A}_{\text{fft}} \mathbf{s}_{\text{fft}} - \frac{q}{p} \left[\frac{p}{q} \mathbf{v}^T \mathbf{A}_{\text{fft}} \right] \tilde{\mathbf{s}}_{\text{fft}} = \mathbf{v}^T \mathbf{A}_{\text{fft}} (\mathbf{s}_{\text{fft}} - \tilde{\mathbf{s}}_{\text{fft}}) + \frac{q}{p} \left\{ \frac{p}{q} \mathbf{v}^T \mathbf{A}_{\text{fft}} \right\} \tilde{\mathbf{s}}_{\text{fft}}.$$

Since $\mathbf{v} \in V$, we have $\mathbf{v}^T \mathbf{A}_{\text{dual}} \mathbf{s}_{\text{dual}} = 0 \pmod{q}$. Hence the whole phase can be written as $\mathbf{v}^T \mathbf{A}_{\text{guess}} \delta + \eta(\mathbf{v}) \pmod{q}$, where

$$\eta(\mathbf{v}) := \mathbf{v}^T \mathbf{A}_{\text{fft}} (\mathbf{s}_{\text{fft}} - \tilde{\mathbf{s}}_{\text{fft}}) + \frac{q}{p} \left\{ \frac{p}{q} \mathbf{v}^T \mathbf{A}_{\text{fft}} \right\} \tilde{\mathbf{s}}_{\text{fft}} + \mathbf{v}^T \mathbf{e}.$$

Thus

$$X_{\mathbf{v}} = \cos \left(\frac{2\pi}{q} (\mathbf{v}^T \mathbf{A}_{\text{guess}} \delta + \eta(\mathbf{v})) \right).$$

Here $\eta(\mathbf{v})$ may depend on $\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}, \mathbf{s}, \mathbf{e}$ and $\tilde{\mathbf{s}}_{\text{fft}}$, but it is independent of $\mathbf{A}_{\text{guess}}$. This is the source of the cancellation below.

Lemma 14. *Let q be prime. Let $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{Z}_q^m$ be arbitrary fixed vectors, and let $\delta = \mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ be nonzero. Denote $\Delta = \mathbf{v}_1 \pm \mathbf{v}_2 \in \mathbb{Z}_q^m$. If $\Delta \neq \mathbf{0} \bmod q$, then for any real-valued random variable C independent of $\mathbf{A}_{\text{guess}}$,*

$$\mathbb{E}_{\mathbf{A}_{\text{guess}}, C} \left[\cos \left(\frac{2\pi}{q} (\Delta^T \mathbf{A}_{\text{guess}} \delta + C) \right) \right] = 0.$$

Proof. The complete proof is given in Appendix I.

By Lemma 14, for fixed $\eta(\mathbf{v})$ and $\mathbf{v} \in V^\times$, if $\delta \neq \mathbf{0}$, then the expectation of $X_{\mathbf{v}}$ over $\mathbf{A}_{\text{guess}}$ is zero. Since the distribution of $\mathbf{w}_j \bmod q$ over $V^\times$ is independent of $\mathbf{A}_{\text{guess}}$, conditioning on the residue class gives

$$\mathbb{E}_{\mathbf{A}, \mathbf{s}, \mathbf{w}_j} [X_j] = \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* \mathbb{E}_{\mathbf{A}_{\text{guess}}} [X_{\mathbf{v}}] = 0.$$

Thus, for any incorrect candidate $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$, and any auxiliary choice $\tilde{\mathbf{s}}_{\text{fft}}$, the individual score has expectation zero. By linearity, $\mathbb{E}[F_{\mathbf{W}}] = \sum_{j=1}^N \mathbb{E}[X_j] = 0$.

Second-Moment Analysis under an Incorrect Guess Denote

$$\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) = \mathbb{E}_{\mathbf{w}_j} [X_j(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) \mid (\mathbf{A}, \mathbf{s}, \mathbf{e})].$$

On the incorrect-candidate side, we need to prove that the total score stays below the threshold

$$F_{\mathbf{W}} < \frac{1}{2} \cdot N D_{\text{fft}} h^*(u_*).$$

We split this threshold into two parts. First, using the second moment and Markov's inequality, we will show that outside a small-probability bad event,

$$|\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})| < \frac{1}{4} D_{\text{fft}} h^*(u_*).$$

Thus, on the corresponding good event,

$$N \mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) \leq N |\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})| < \frac{1}{4} \cdot N D_{\text{fft}} h^*(u_*).$$

Second, on this good event, conditional Hoeffding will control the deviation from the conditional total mean:

$$F_{\mathbf{W}} - N \mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) < \frac{1}{4} \cdot N D_{\text{fft}} h^*(u_*).$$

Combining the two inequalities gives $F_{\mathbf{W}} < \frac{1}{2}ND_{\text{fft}}h^*(u_*)$.

It remains in this subsection to compute $\mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})}[\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})^2]$, which will be used with Theorem 3. By the law of total expectation,

$$\begin{aligned} & \mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} [\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})^2] \\ &= \mathbb{E}_{(\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}, \mathbf{s}, \mathbf{e})} \left[\sum_{\mathbf{v}_1, \mathbf{v}_2 \in V^\times} p_{\mathbf{v}_1}^* p_{\mathbf{v}_2}^* \mathbb{E}_{\mathbf{A}_{\text{guess}}} [X_{\mathbf{v}_1} X_{\mathbf{v}_2} \mid (\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}, \mathbf{s}, \mathbf{e})] \right]. \end{aligned}$$

Next, we estimate $\mathbb{E}_{\mathbf{A}_{\text{guess}}} [X_{\mathbf{v}_1} X_{\mathbf{v}_2} \mid (\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}, \mathbf{s}, \mathbf{e})]$ according to the relation between $\mathbf{v}_1, \mathbf{v}_2 \in V^\times$.

Case 1: Diagonal term $\mathbf{v}_1 = \mathbf{v}_2 = \mathbf{v}$. Here $\mathbf{v}_1 - \mathbf{v}_2 = \mathbf{0}$, while $\mathbf{v}_1 + \mathbf{v}_2 = 2\mathbf{v} \neq \mathbf{0}$, since q is an odd prime. The product-to-sum formula gives

$$X_{\mathbf{v}}^2 = \frac{1}{2} + \frac{1}{2} \cos \left(\frac{2\pi}{q} (2\mathbf{v}^T \mathbf{A}_{\text{guess}} \delta + 2\eta(\mathbf{v})) \right).$$

In the second term, $\Delta = 2\mathbf{v} \neq \mathbf{0}$. Hence Lemma 14 implies that its expectation over $\mathbf{A}_{\text{guess}}$ is zero, and therefore

$$\mathbb{E}_{\mathbf{A}_{\text{guess}}} [X_{\mathbf{v}}^2 \mid (\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}, \mathbf{s}, \mathbf{e})] = \frac{1}{2}.$$

Case 2: Anti-diagonal term $\mathbf{v}_1 = -\mathbf{v}_2$. Let $\mathbf{v}_2 = -\mathbf{v}_1$, with $\mathbf{v}_1 \in V^\times$. Then $\mathbf{v}_1 + \mathbf{v}_2 = \mathbf{0}$, while $\mathbf{v}_1 - \mathbf{v}_2 = 2\mathbf{v}_1 \neq \mathbf{0}$. The product-to-sum formula gives

$$\begin{aligned} X_{\mathbf{v}_1} X_{-\mathbf{v}_1} &= \frac{1}{2} \cos \left(\frac{2\pi}{q} (\eta(\mathbf{v}_1) + \eta(-\mathbf{v}_1)) \right) \\ &\quad + \frac{1}{2} \cos \left(\frac{2\pi}{q} (2\mathbf{v}_1^T \mathbf{A}_{\text{guess}} \delta + \eta(\mathbf{v}_1) - \eta(-\mathbf{v}_1)) \right). \end{aligned}$$

The second term has zero expectation over $\mathbf{A}_{\text{guess}}$ by Lemma 14. Hence

$$\mathbb{E}_{\mathbf{A}_{\text{guess}}} [X_{\mathbf{v}_1} X_{-\mathbf{v}_1} \mid (\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}, \mathbf{s}, \mathbf{e})] = \frac{1}{2} \cos \left(\frac{2\pi}{q} (\eta(\mathbf{v}_1) + \eta(-\mathbf{v}_1)) \right),$$

and therefore

$$|\mathbb{E}_{\mathbf{A}_{\text{guess}}} [X_{\mathbf{v}_1} X_{-\mathbf{v}_1} \mid (\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}, \mathbf{s}, \mathbf{e})]| \leq \frac{1}{2}.$$

Case 3: Remaining terms $\mathbf{v}_1 \neq \pm \mathbf{v}_2$. In this case $\mathbf{v}_1 + \mathbf{v}_2 \neq \mathbf{0}$ and $\mathbf{v}_1 - \mathbf{v}_2 \neq \mathbf{0}$. After applying the product-to-sum formula, both terms have zero expectation over $\mathbf{A}_{\text{guess}}$ by Lemma 14. Thus

$$\mathbb{E}_{\mathbf{A}_{\text{guess}}} [X_{\mathbf{v}_1} X_{\mathbf{v}_2} \mid (\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}, \mathbf{s}, \mathbf{e})] = 0.$$

Combining the three cases, we obtain the uniform bound

$$|\mathbb{E}_{\mathbf{A}_{\text{guess}}} [X_{\mathbf{v}_1} X_{\mathbf{v}_2} \mid (\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}, \mathbf{s}, \mathbf{e})]| \leq \begin{cases} 0, & \mathbf{v}_1 \neq \pm \mathbf{v}_2, \\ \frac{1}{2}, & \text{otherwise.} \end{cases}$$

Substituting this case-by-case estimate into the residue-class expansion of the second moment gives

$$\mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} \left[\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})^2 \right] \leq \frac{1}{2} \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^{*2} + \frac{1}{2} \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* p_{-\mathbf{v}}^*.$$

Using

$$p_{\mathbf{v}}^* p_{-\mathbf{v}}^* \leq \frac{1}{2} (p_{\mathbf{v}}^{*2} + p_{-\mathbf{v}}^{*2}),$$

we have

$$\sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* p_{-\mathbf{v}}^* \leq \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^{*2}.$$

Therefore

$$\mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} \left[\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})^2 \right] \leq \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^{*2} \leq p_{\max}^* \leq \bar{p}_{\max}^*.$$

Conditional Hoeffding Inequality and Separation In this subsection, we first bound the probability that the conditional mean of an incorrect candidate is too large. By the second-moment estimate from the previous subsection and Theorem 3, we obtain

$$\begin{aligned} & \Pr_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} \left(|\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})| \geq \frac{1}{4} D_{\text{fft}} h^*(u_*) \right) \\ & \leq \frac{\mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} \left[\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})^2 \right]}{\left(\frac{1}{4} D_{\text{fft}} h^*(u_*) \right)^2} \\ & \leq \frac{16 \bar{p}_{\max}^*}{(D_{\text{fft}} h^*(u_*))^2}. \end{aligned} \tag{11}$$

Define the good event

$$\mathcal{G} := \left\{ |\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})| < \frac{1}{4} D_{\text{fft}} h^*(u_*) \right\}.$$

On $\mathcal{G}$, the conditional mean of the individual incorrect score is bounded by $\frac{1}{4} D_{\text{fft}} h^*(u_*)$. On the bad event $\mathcal{G}^c$, we do not control this conditional mean; instead, we only use (11) to bound the probability of $\mathcal{G}^c$.

We next bound the probability $p_{\text{err}} := \Pr(F_{\mathbf{W}} \geq \frac{1}{2} N D_{\text{fft}} h^*(u_*))$ for a fixed incorrect candidate. By the law of total probability, conditioning on the complete instance gives

$$p_{\text{err}} = \mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} \left[\Pr \left(F_{\mathbf{W}} \geq \frac{1}{2} N D_{\text{fft}} h^*(u_*) \mid (\mathbf{A}, \mathbf{s}, \mathbf{e}) \right) \right].$$

Splitting according to $\mathcal{G}$ and $\mathcal{G}^c$, we have

$$p_{\text{err}} = \mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} \left[\mathbf{1}_{\mathcal{G}} \Pr \left(F_{\mathbf{W}} \geq \frac{1}{2} N D_{\text{fft}} h^*(u_*) \mid (\mathbf{A}, \mathbf{s}, \mathbf{e}) \right) \right] \\ + \mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} \left[\mathbf{1}_{\mathcal{G}^c} \Pr \left(F_{\mathbf{W}} \geq \frac{1}{2} N D_{\text{fft}} h^*(u_*) \mid (\mathbf{A}, \mathbf{s}, \mathbf{e}) \right) \right]. \quad (12)$$

Here $\mathbf{1}_{\mathcal{G}} := \mathbf{1}_{\mathcal{G}}(\mathbf{A}, \mathbf{s}, \mathbf{e})$ is the indicator of $\mathcal{G}$.

Conditional probability on the good event $\mathcal{G}$. Fix $(\mathbf{A}, \mathbf{s}, \mathbf{e}) \in \mathcal{G}$. Conditioned on this instance, the short vectors $\mathbf{w}_1, \dots, \mathbf{w}_N$ are independent and identically distributed, and $X_j \in [-1, 1]$. Moreover,

$$\mathbb{E}_{\mathbf{w}_j} [X_j (\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) \mid (\mathbf{A}, \mathbf{s}, \mathbf{e})] = \mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}),$$

and on $\mathcal{G}$,

$$\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) \leq |\mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}})| < \frac{1}{4} D_{\text{fft}} h^*(u_*).$$

Therefore,

$$F_{\mathbf{W}} \geq \frac{1}{2} N D_{\text{fft}} h^*(u_*) \implies F_{\mathbf{W}} - N \mu(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) \geq \frac{1}{4} N D_{\text{fft}} h^*(u_*).$$

Applying Theorem 1, we get

$$\Pr \left(F_{\mathbf{W}} \geq \frac{1}{2} N D_{\text{fft}} h^*(u_*) \mid (\mathbf{A}, \mathbf{s}, \mathbf{e}) \right) \leq \exp \left(-\frac{N(D_{\text{fft}} h^*(u_*))^2}{32} \right).$$

Trivial upper bound on the bad event $\mathcal{G}^c$. For $(\mathbf{A}, \mathbf{s}, \mathbf{e}) \in \mathcal{G}^c$, we use only the trivial bound

$$\Pr \left(F_{\mathbf{W}} \geq \frac{1}{2} N D_{\text{fft}} h^*(u_*) \mid (\mathbf{A}, \mathbf{s}, \mathbf{e}) \right) \leq 1.$$

Combining the two terms. Substituting the two bounds into (12), we obtain

$$p_{\text{err}} \leq \mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} \left[\mathbf{1}_{\mathcal{G}} \exp \left(-\frac{N(D_{\text{fft}} h^*(u_*))^2}{32} \right) \right] + \mathbb{E}_{(\mathbf{A}, \mathbf{s}, \mathbf{e})} [\mathbf{1}_{\mathcal{G}^c}] \\ = \exp \left(-\frac{N(D_{\text{fft}} h^*(u_*))^2}{32} \right) \Pr(\mathcal{G}) + \Pr(\mathcal{G}^c) \\ \leq \exp \left(-\frac{N(D_{\text{fft}} h^*(u_*))^2}{32} \right) + \frac{16\bar{p}_{\text{max}}^*}{(D_{\text{fft}} h^*(u_*))^2}.$$

Overall success probability. Denote $p_{\text{cor}}^{\min} := p_{\text{cor}} \left((h^*(u_*))^2 \right)$. Combining the correct-candidate analysis, the incorrect-candidate false-positive bound, and the full-rank event for $\mathbf{A}_{\text{dual}}$, we obtain

$$p_{\text{success}} \geq 1 - (1 - p_{\text{cor}}^{\min}) - (2\eta + 1)^{n_{\text{guess}} + n_{\text{fft}}} p_{\text{err}} - (1 - p_{\text{rank}}) \\ = p_{\text{cor}}^{\min} - (2\eta + 1)^{n_{\text{guess}} + n_{\text{fft}}} p_{\text{err}} + p_{\text{rank}} - 1.$$

When the parameters (m, n, q, s) satisfy

$$\bar{p}_{\max}^* \leq \frac{1}{100} \cdot \frac{1}{(2\eta + 1)^{n_{\text{guess}} + n_{\text{fft}}}} \cdot \frac{1}{N^2},$$

the covariance tail term $N(N-1)(q-1)\bar{p}_{\max}^*$ is sufficiently small for the above probability bound. We call the range of parameters for which $p_{\text{success}} > 0.01$ the success region of the dual attack. This definition is based on the conditional-covariance analysis, rather than only on the geometric properties of random lattices as in [24, 26].

Combining the correct-candidate lower bound with the false-positive bound for all incorrect candidates, we obtain the following separation statement. With probability at least p_{success} , the true candidate satisfies

$$F_{\mathbf{W}}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) > \frac{1}{2}ND_{\text{fft}}h^*(u_*),$$

whereas every candidate with $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$ satisfies

$$F_{\mathbf{W}}(\tilde{\mathbf{s}}_{\text{guess}}, \tilde{\mathbf{s}}_{\text{fft}}) < \frac{1}{2}ND_{\text{fft}}h^*(u_*).$$

Therefore, thresholding the FFT score at $\frac{1}{2}ND_{\text{fft}}h^*(u_*)$ recovers $\mathbf{s}_{\text{guess}}$.

This completes the correctness analysis. The correct-candidate side and the incorrect-candidate side use different conditioning orders because they serve different purposes: the former proves a lower bound for the true score, while the latter proves an upper bound for all wrong recovered guesses.

4 Attack Summary and Informal Cost Estimation

We summarize the preceding correctness analysis as the following theorem. The theorem gives the correctness criterion used to determine the required number of short vectors.

To keep the comparison meaningful, we deliberately use the same informal cost-estimation convention as [23] for all shared subroutines; only this correctness criterion is changed.

Theorem 8 (Correctness and cost of the MATZOV-style FFT dual attack). *Let $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q)$ be an LWE instance from the ML-KEM parameter sets, where*

$$\mathbf{A} = (\mathbf{A}_{\text{guess}}, \mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{dual}}), \quad \mathbf{s}^T = (\mathbf{s}_{\text{guess}}^T, \mathbf{s}_{\text{fft}}^T, \mathbf{s}_{\text{dual}}^T).$$

Let $\Lambda = L_q^\perp(\mathbf{A}_{\text{dual}})$, and let $\mathbf{w}_1, \dots, \mathbf{w}_N$ be independently sampled from the conditional discrete Gaussian distribution over Λ defined in Section 2.6.

Let $u_* := \sqrt{m\eta}/2$ be the error-norm estimate used in the concrete analysis, and define

$$h^*(u_*) := \frac{\rho_{1/s}(u_*) - \frac{B}{|V|}}{1 - \frac{B}{|V|}}.$$

We assume $B < |V|$, $B/|V| < \rho_{1/s}(u_*)$, and $D_{\text{fft}} > 0$, so that $\bar{p}_{\text{max}}^*$ is well-defined and $D_{\text{fft}}h^*(u_*) > 0$. Choose

$$N := \left\lceil \frac{32(n_{\text{guess}} + n_{\text{fft}}) \ln(2\eta + 1) + 32 \ln(100)}{(D_{\text{fft}}h^*(u_*))^2} \right\rceil.$$

Then accepting candidates whose FFT score is at least $C := ND_{\text{fft}}h^*(u_*)/2$ recovers $\mathbf{s}_{\text{guess}}$ with probability at least

$$p_{\text{success}} \geq p_{\text{cor}}^{\min} - (2\eta + 1)^{n_{\text{guess}} + n_{\text{fft}}} p_{\text{err}} + p_{\text{rank}} - 1.$$

The one-run time complexity is upper bounded by

$$\begin{aligned} T_{\text{one}} &\leq T_{\text{BKZ}}(m, \beta) + \frac{N}{1 - B/|V|} \cdot T_{\text{MCMC}}(L_q^\perp(\mathbf{A}_{\text{dual}}), qs) \\ &\quad + O(N + p^{n_{\text{fft}}})(2\eta + 1)^{n_{\text{guess}}}. \end{aligned}$$

and the reported expected cost is

$$T_{\text{total}} = \frac{T_{\text{one}}}{p_{\text{success}}}.$$

We instantiate this theorem for ML-KEM. To ensure a meaningful comparison with previous provable dual attacks, we adopt the same cost-estimation conventions as in [23,25,26] for identical subroutines. Specifically, BKZ and MCMC costs follow [23], while our analysis determines the required number N of short vectors. Under the same corrected σ_e convention, [26] reports 243, 358, and 495 bits, which are higher than those in [23]; consequently, Table 2 utilizes [23] as the main baseline.

Table 2. Concrete cost estimates for ML-KEM parameters. Here “Red.” denotes the bit reduction relative to [23].

Scheme	[23]	Ours	Red.	m	n_{guess}	n_{fft}	n_{dual}	p_{success}	p	β	s
ML-KEM-512	238	210	28	1044	24	58	430	0.0337	5	620	0.084
ML-KEM-768	347	300	47	1544	53	74	641	0.0853	5	957	0.097
ML-KEM-1024	478	410	68	2048	99	75	850	0.0839	5	1343	0.091

The results are reported in Table 2. Compared with the corrected estimates of [23], our estimates are reduced by 28, 47, and 68 bits for ML-KEM-512, ML-KEM-768, and ML-KEM-1024, respectively. The improvement becomes more pronounced for larger parameter sets, reflecting the increasing impact of the FFT-accelerated guessing phase.

References

1. Aharonov, D., Regev, O.: Lattice problems in $\text{np} \cap \text{comp. J. ACM}$ **52**(5), 749765 (Sep 2005). <https://doi.org/10.1145/1089023.1089025>, <https://doi.org/10.1145/1089023.1089025>
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-Hilfinger, C., Shen, Y., Tillich, J.P.: Assessing the impact of a variant of matzovs 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. Cheviguard, 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, In pulles. *Journal of Cryptology* **39**(1), 8 (2026)
12. Glaser, T., May, A.: How to enumerate lwe keys as narrow as in kyber/dilithium. In: *International Conference on Cryptology and Network Security*. pp. 75–100. Springer (2023)
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. Leveil, 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. Li, Y., Zheng, Q.: Unified dual attack analyses: Covariance-based score distribution prediction for LWE. Cryptology ePrint Archive, Paper 2026/1048 (2026), <https://eprint.iacr.org/2026/1048>
19. Li, Y., Zheng, Q.: What happens when integrating modulus switching and lossy source coding: A new dual attack variant on LWE. Cryptology ePrint Archive, Paper 2026/1400 (2026), <https://eprint.iacr.org/2026/1400>
20. MATZOV: Report on the security of lwe: Improved dual lattice attack. Tech. rep., IACR Cryptol. ePrint Arch. (2022), report 2022/445
21. 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>
22. Ogilvie, T.: On the concrete hardness gap between MLWE and LWE. Cryptology ePrint Archive, Paper 2026/279 (2026), <https://eprint.iacr.org/2026/279>
23. Pouly, A., Shen, Y.: Provable dual attacks on learning with errors. Cryptology ePrint Archive, Paper 2023/1508 (2023). https://doi.org/10.1007/978-3-031-58754-2_10, <https://eprint.iacr.org/2023/1508>
24. 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)
25. 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)
26. 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)
27. 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>
28. Serre, J.P.: A Course in Arithmetic. Springer (1973)
29. 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>
30. Wang, R.J., Yang, H.S., Wang, Z.X., Zheng, Q.X., Li, Y.C.: Too far behind? narrowing the gap with a dual-enhanced two-stage algebraic algorithm for LWE. Cryptology ePrint Archive, Paper 2026/688 (2026), <https://eprint.iacr.org/2026/688>
31. 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>
32. 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

Proof. We prove a slightly stronger statement than the one needed in the main text. The point is that, once a nonzero residue class of the short vector is fixed, the projections of the independent uniform columns of $\mathbf{A}_{\text{guess}}$ and $\mathbf{A}_{\text{fft}}$ remain uniform. Since the resulting conditional distribution does not depend on the fixed value of $\mathbf{w}$, this also gives independence from $\mathbf{w}$ itself.

Fix any nonzero vector $\mathbf{w} \in \mathbb{Z}_q^m$. Since q is prime and $\mathbf{w} \neq \mathbf{0}$, at least one coordinate of $\mathbf{w}$ is nonzero. Hence the linear map

$$\varphi_{\mathbf{w}} : \mathbb{Z}_q^m \rightarrow \mathbb{Z}_q, \quad \mathbf{x} \mapsto \mathbf{w}^T \mathbf{x}$$

is a nonzero linear functional over the field $\mathbb{Z}_q$. In particular, it is surjective. Therefore, for every $c \in \mathbb{Z}_q$, the affine hyperplane

$$\{\mathbf{x} \in \mathbb{Z}_q^m : \mathbf{w}^T \mathbf{x} = c\}$$

has cardinality q^{m-1} .

Now fix

$$(\mathbf{u}, \mathbf{v}) \in \mathbb{Z}_q^{n_{\text{guess}}} \times \mathbb{Z}_q^{n_{\text{fft}}}.$$

Conditioned on this fixed value of $\mathbf{w}$, the columns of $\mathbf{A}_{\text{guess}}$ and $\mathbf{A}_{\text{fft}}$ are still independent and uniform over $\mathbb{Z}_q^m$. For each column of $\mathbf{A}_{\text{guess}}$, the constraint that its projection under $\varphi_{\mathbf{w}}$ equals the corresponding coordinate of $\mathbf{u}$ leaves exactly q^{m-1} possible columns out of q^m . The same argument applies independently to each column of $\mathbf{A}_{\text{fft}}$. Hence

$$\Pr[\mathbf{w}^T \mathbf{A}_{\text{guess}} = \mathbf{u}^T, \mathbf{w}^T \mathbf{A}_{\text{fft}} = \mathbf{v}^T \mid \mathbf{w}] = \left(\frac{q^{m-1}}{q^m} \right)^{n_{\text{guess}} + n_{\text{fft}}} = q^{-(n_{\text{guess}} + n_{\text{fft}})}.$$

Thus, conditioned on $\mathbf{w}$, the pair

$$(\mathbf{w}^T \mathbf{A}_{\text{guess}}, \mathbf{w}^T \mathbf{A}_{\text{fft}})$$

is uniform over $\mathbb{Z}_q^{n_{\text{guess}}} \times \mathbb{Z}_q^{n_{\text{fft}}}$.

Moreover, the above conditional probability does not depend on the particular nonzero value of $\mathbf{w}$. Therefore, for every nonzero $\mathbf{w}$ and every $(\mathbf{u}, \mathbf{v})$,

$$\Pr[\mathbf{w}^T \mathbf{A}_{\text{guess}} = \mathbf{u}^T, \mathbf{w}^T \mathbf{A}_{\text{fft}} = \mathbf{v}^T \mid \mathbf{w}] = \Pr[\mathbf{w}^T \mathbf{A}_{\text{guess}} = \mathbf{u}^T, \mathbf{w}^T \mathbf{A}_{\text{fft}} = \mathbf{v}^T].$$

This proves that the pair

$$(\mathbf{w}^T \mathbf{A}_{\text{guess}}, \mathbf{w}^T \mathbf{A}_{\text{fft}})$$

is independent of $\mathbf{w}$.

Finally, the same displayed probability factors as

$$q^{-(n_{\text{guess}} + n_{\text{fft}})} = q^{-n_{\text{guess}}} q^{-n_{\text{fft}}},$$

which shows that $\mathbf{w}^T \mathbf{A}_{\text{guess}}$ and $\mathbf{w}^T \mathbf{A}_{\text{fft}}$ are independent. Combining this block independence with the independence from $\mathbf{w}$, we obtain the claimed mutual independence of $\mathbf{w}, \mathbf{w}^T \mathbf{A}_{\text{guess}}, \mathbf{w}^T \mathbf{A}_{\text{fft}}$.

B Proof of Lemma 7

Proof. The proof has two steps. First we compute the Fourier expectation before discarding the residue class $\mathbf{0} \bmod q$. Then we remove the contribution of that residue class and renormalize by the acceptance probability $1 - p_0$.

Without conditioning on $\mathbf{w} \bmod q \neq \mathbf{0}$, the Poisson summation formula gives

$$\mathbb{E}_{\mathbf{w} \leftarrow D_{\Lambda, qs}} \left[e^{2\pi i \langle \mathbf{w}, \mathbf{e}/q \rangle} \right] = \frac{\rho_{1/(qs)}(\mathbf{e}/q + \widehat{\Lambda})}{\rho_{1/(qs)}(\widehat{\Lambda})} = \frac{\rho_{1/s}(\mathbf{e} + q\widehat{\Lambda})}{\rho_{1/s}(q\widehat{\Lambda})} = f_{L, 1/s}(\mathbf{e}),$$

where $L = q\widehat{\Lambda}$. Here the second equality follows by scaling the lattice and the shift by q .

We now remove the residue class $\mathbf{w} \bmod q = \mathbf{0}$. Let

$$E_0 := \{\mathbf{w} \bmod q = \mathbf{0}\}.$$

By definition, $\Pr[E_0 \mid \mathbf{A}_{\text{dual}}] = p_0$. Moreover, on E_0 , we have $\mathbf{w} = q\mathbf{z}$ for some $\mathbf{z} \in \mathbb{Z}^m$. Since $\mathbf{e} \in \mathbb{Z}^m$, this gives

$$e^{2\pi i \mathbf{w}^T \mathbf{e}/q} = e^{2\pi i \mathbf{z}^T \mathbf{e}} = 1.$$

Therefore the contribution of E_0 to the unconditioned expectation is exactly p_0 . Splitting the expectation according to E_0 and E_0^c , we get

$$f_{L, 1/s}(\mathbf{e}) = p_0 + (1 - p_0) \mathbb{E}_{\mathbf{w}_j} \left[e^{\frac{2\pi i}{q} \mathbf{w}_j^T \mathbf{e}} \mid (\mathbf{A}_{\text{dual}}, \mathbf{e}) \right],$$

where the expectation on the right is taken over $D_{\Lambda, qs}$ conditioned on $\mathbf{w}_j \bmod q \neq \mathbf{0}$. Rearranging gives

$$\mathbb{E}_{\mathbf{w}_j} \left[e^{\frac{2\pi i}{q} \mathbf{w}_j^T \mathbf{e}} \mid (\mathbf{A}_{\text{dual}}, \mathbf{e}) \right] = \frac{f_{L, 1/s}(\mathbf{e}) - p_0}{1 - p_0}.$$

On the other hand, conditioning further on the nonzero residue class $\mathbf{w}_j \bmod q = \mathbf{v}$ gives

$$\mathbb{E}_{\mathbf{w}_j} \left[e^{\frac{2\pi i}{q} \mathbf{w}_j^T \mathbf{e}} \mid (\mathbf{A}_{\text{dual}}, \mathbf{e}) \right] = \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* e^{\frac{2\pi i}{q} \mathbf{v}^T \mathbf{e}} = \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* e^{i\theta_{\mathbf{v}}}.$$

This is the form used in the covariance computation: after averaging over the short vector, its randomness appears only through the normalized weights $p_{\mathbf{v}}^*$ on nonzero residue classes.

Finally, since the discrete Gaussian is symmetric under $\mathbf{w} \mapsto -\mathbf{w}$, we have $p_{\mathbf{v}}^* = p_{-\mathbf{v}}^*$. Together with $\theta_{-\mathbf{v}} = -\theta_{\mathbf{v}}$, this shows that the imaginary parts cancel in pairs. Hence the above quantity is real.

C Proof of Lemma 8

Proof. By Lemma 7, we have

$$h^*(\mathbf{e}) = \sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* e^{i\theta_{\mathbf{v}}},$$

and this quantity is real. We now expand the two double sums directly.

By Euler's formula,

$$\cos(\theta_{\mathbf{v}} + \theta_{\mathbf{w}}) = \Re \left[e^{i(\theta_{\mathbf{v}} + \theta_{\mathbf{w}})} \right] = \Re \left[e^{i\theta_{\mathbf{v}}} e^{i\theta_{\mathbf{w}}} \right].$$

Therefore,

$$\begin{aligned} \sum_{\mathbf{v}, \mathbf{w} \in V^\times} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} + \theta_{\mathbf{w}}) &= \Re \left[\sum_{\mathbf{v}, \mathbf{w} \in V^\times} p_{\mathbf{v}}^* p_{\mathbf{w}}^* e^{i\theta_{\mathbf{v}}} e^{i\theta_{\mathbf{w}}} \right] \\ &= \Re \left[\left(\sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* e^{i\theta_{\mathbf{v}}} \right) \left(\sum_{\mathbf{w} \in V^\times} p_{\mathbf{w}}^* e^{i\theta_{\mathbf{w}}} \right) \right] = \Re \left[h^*(\mathbf{e})^2 \right]. \end{aligned}$$

Since $h^*(\mathbf{e}) \in \mathbb{R}$, the last expression equals $(h^*(\mathbf{e}))^2$.

We next prove the identity with the minus sign. Again by Euler's formula,

$$\cos(\theta_{\mathbf{v}} - \theta_{\mathbf{w}}) = \Re \left[e^{i(\theta_{\mathbf{v}} - \theta_{\mathbf{w}})} \right] = \Re \left[e^{i\theta_{\mathbf{v}}} e^{-i\theta_{\mathbf{w}}} \right].$$

Thus,

$$\begin{aligned} \sum_{\mathbf{v}, \mathbf{w} \in V^\times} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} - \theta_{\mathbf{w}}) &= \Re \left[\sum_{\mathbf{v}, \mathbf{w} \in V^\times} p_{\mathbf{v}}^* p_{\mathbf{w}}^* e^{i\theta_{\mathbf{v}}} e^{-i\theta_{\mathbf{w}}} \right] \\ &= \Re \left[\left(\sum_{\mathbf{v} \in V^\times} p_{\mathbf{v}}^* e^{i\theta_{\mathbf{v}}} \right) \left(\sum_{\mathbf{w} \in V^\times} p_{\mathbf{w}}^* e^{-i\theta_{\mathbf{w}}} \right) \right]. \end{aligned}$$

Since $p_{\mathbf{w}}^* \in \mathbb{R}$, we have

$$\sum_{\mathbf{w} \in V^\times} p_{\mathbf{w}}^* e^{-i\theta_{\mathbf{w}}} = \overline{\sum_{\mathbf{w} \in V^\times} p_{\mathbf{w}}^* e^{i\theta_{\mathbf{w}}}} = \overline{h^*(\mathbf{e})}.$$

Consequently,

$$\sum_{\mathbf{v}, \mathbf{w} \in V^\times} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} - \theta_{\mathbf{w}}) = \Re \left[h^*(\mathbf{e}) \overline{h^*(\mathbf{e})} \right] = |h^*(\mathbf{e})|^2.$$

Finally, since $h^*(\mathbf{e}) \in \mathbb{R}$, we have

$$|h^*(\mathbf{e})|^2 = (h^*(\mathbf{e}))^2.$$

This proves both identities.

D Proof of Lemma 10

Proof. We first prove the uniform modulus bound for the single-column quantities H^+ and H^- . Since $|g_{s_t}(x)| = 1$ for every $x \in \mathbb{Z}_q$, we have

$$|g_{s_t}(\mathbf{v}^\top \mathbf{a}) g_{s_t}(\mathbf{w}^\top \mathbf{a})| = 1, \quad \left| g_{s_t}(\mathbf{v}^\top \mathbf{a}) \overline{g_{s_t}(\mathbf{w}^\top \mathbf{a})} \right| = 1.$$

Taking expectations and using the triangle inequality gives

$$|H^+(\mathbf{v}, \mathbf{w})| \leq 1, \quad |H^-(\mathbf{v}, \mathbf{w})| \leq 1.$$

We next pass from the single-column quantities to the full products $\alpha_{\mathbf{v}}$ and $\alpha_{\mathbf{w}}$. By definition,

$$\alpha_{\mathbf{v}} = \prod_{t=1}^{n_{\text{fft}}} g_{s_t}(\mathbf{v}^\top \mathbf{a}_t), \quad \alpha_{\mathbf{w}} = \prod_{t=1}^{n_{\text{fft}}} g_{s_t}(\mathbf{w}^\top \mathbf{a}_t),$$

where $\mathbf{a}_t$ is the t -th column of $\mathbf{A}_{\text{fft}}$. The pairs $(\mathbf{a}_t, s_t)$ are independent over different values of t , and each pair has the same distribution. Therefore the expectation factors column by column:

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \alpha_{\mathbf{w}}] = \prod_{t=1}^{n_{\text{fft}}} \mathbb{E}_{\mathbf{a}_t, s_t} [g_{s_t}(\mathbf{v}^\top \mathbf{a}_t) g_{s_t}(\mathbf{w}^\top \mathbf{a}_t)] = H^+(\mathbf{v}, \mathbf{w})^{n_{\text{fft}}},$$

and similarly

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \overline{\alpha_{\mathbf{w}}}] = \prod_{t=1}^{n_{\text{fft}}} \mathbb{E}_{\mathbf{a}_t, s_t} [g_{s_t}(\mathbf{v}^\top \mathbf{a}_t) \overline{g_{s_t}(\mathbf{w}^\top \mathbf{a}_t)}] = H^-(\mathbf{v}, \mathbf{w})^{n_{\text{fft}}}.$$

It remains to specialize these identities to linearly independent pairs. If $(\mathbf{v}, \mathbf{w}) \in \mathcal{I}$, Lemma 9 gives the single-column values

$$H^+(\mathbf{v}, \mathbf{w}) = \mathbb{E}_{s_t} [c_0(s_t)^2], \\ H^-(\mathbf{v}, \mathbf{w}) = \mathbb{E}_{s_t} [|c_0(s_t)|^2].$$

Substituting these two identities into the factored expressions above yields

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \alpha_{\mathbf{w}}] = (\mathbb{E}_{s_t} [c_0(s_t)^2])^{n_{\text{fft}}} = \kappa_{\text{sq}},$$

and

$$\mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}} [\alpha_{\mathbf{v}} \overline{\alpha_{\mathbf{w}}}] = (\mathbb{E}_{s_t} [|c_0(s_t)|^2])^{n_{\text{fft}}} = \kappa_{\text{abs}}.$$

This proves all the claimed identities.

E Proof of Theorem 6

Proof. We give the details of the estimate. The only loss comes from replacing the sums over linearly independent residue-class pairs by the corresponding sums over all pairs, and then bounding the omitted dependent-pair contribution.

The identity (9) follows directly from the covariance decomposition, Lemma 11, and Theorem 5.

It remains to prove the stated bound. Recall that $\mathcal{D}$ contains the linearly dependent pairs, while $\mathcal{I}$ contains the linearly independent pairs. Define

$$S_+^{\mathcal{D}} := \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{D}} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} + \theta_{\mathbf{w}}),$$

and

$$S_-^{\mathcal{D}} := \sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{D}} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} - \theta_{\mathbf{w}}).$$

By Lemma 8, the two corresponding cosine sums over all pairs in $(V^\times)^2$ both equal $(h^*(\mathbf{e}))^2$. Since $(V^\times)^2$ is the disjoint union of $\mathcal{I}$ and $\mathcal{D}$, the sums over $\mathcal{I}$ are obtained by subtracting the dependent-pair parts:

$$\sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{I}} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} + \theta_{\mathbf{w}}) = (h^*(\mathbf{e}))^2 - S_+^{\mathcal{D}},$$

and

$$\sum_{(\mathbf{v}, \mathbf{w}) \in \mathcal{I}} p_{\mathbf{v}}^* p_{\mathbf{w}}^* \cos(\theta_{\mathbf{v}} - \theta_{\mathbf{w}}) = (h^*(\mathbf{e}))^2 - S_-^{\mathcal{D}}.$$

Substituting these two identities into the definition of $Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})$ gives

$$Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) = \frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} (h^*(\mathbf{e}))^2 - \frac{\kappa_{\text{sq}}}{2} S_+^{\mathcal{D}} - \frac{\kappa_{\text{abs}}}{2} S_-^{\mathcal{D}}.$$

Moreover, by Lemma 11,

$$(\bar{\mu}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}))^2 = D_{\text{fft}}^2 (h^*(\mathbf{e}))^2.$$

Therefore

$$\begin{aligned} & Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) - (\bar{\mu}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}))^2 \\ &= \left(\frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} - D_{\text{fft}}^2 \right) (h^*(\mathbf{e}))^2 - \frac{\kappa_{\text{sq}}}{2} S_+^{\mathcal{D}} - \frac{\kappa_{\text{abs}}}{2} S_-^{\mathcal{D}}. \end{aligned}$$

The last two terms are supported only on $\mathcal{D}$. Since $|\cos(\theta_{\mathbf{v}} \pm \theta_{\mathbf{w}})| \leq 1$, we have

$$|S_+^{\mathcal{D}}| \leq W_{\mathcal{D}}, \quad |S_-^{\mathcal{D}}| \leq W_{\mathcal{D}}.$$

Also, Lemma 10 implies $|\kappa_{\text{sq}}| \leq 1$ and $|\kappa_{\text{abs}}| \leq 1$. Hence

$$\left| Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) - (\bar{\mu}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}))^2 - \left(\frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} - D_{\text{fft}}^2 \right) (h^*(\mathbf{e}))^2 \right| \leq W_{\mathcal{D}}.$$

On the other hand, Theorem 5 gives

$$|\varepsilon(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})| \leq W_{\mathcal{D}}.$$

Combining this with

$$\text{Cov} = Q(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}) - (\bar{\mu}(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}}))^2 + \varepsilon(\mathbf{s}_{\text{guess}}, \mathbf{s}_{\text{fft}})$$

and applying the triangle inequality, we obtain

$$\left| \text{Cov} - \left(\frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} - D_{\text{fft}}^2 \right) (h^*(\mathbf{e}))^2 \right| \leq 2W_{\mathcal{D}}.$$

Finally, Lemma 6 gives

$$W_{\mathcal{D}} \leq (q-1)p_{\max}^*,$$

and therefore

$$\left| \text{Cov} - \left(\frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} - D_{\text{fft}}^2 \right) (h^*(\mathbf{e}))^2 \right| \leq 2(q-1)p_{\max}^*.$$

This proves the claimed bound.

F Proof of Lemma 12

Proof. We prove the bound by estimating the Gaussian mass of each residue class modulo q inside the dual lattice.

Since $\mathbf{A}_{\text{dual}}$ has full column rank, we have $|V| = q^{m-n_{\text{dual}}}$ and $\det(\Lambda) = q^{n_{\text{dual}}}$. Fix $\mathbf{v} \in V$. By the definition of V , there exists a lift $\mathbf{x}_{\mathbf{v}} \in \Lambda$ such that $\mathbf{x}_{\mathbf{v}} \bmod q = \mathbf{v}$. Moreover, since $q\mathbb{Z}^m \subseteq \Lambda$, the elements of Λ with residue class $\mathbf{v}$ modulo q are exactly

$$\Lambda_{\mathbf{v}} := \{\mathbf{x} \in \Lambda : \mathbf{x} \bmod q = \mathbf{v}\} = \mathbf{x}_{\mathbf{v}} + q\mathbb{Z}^m.$$

Indeed, any two elements of $\Lambda_{\mathbf{v}}$ differ by a vector in $q\mathbb{Z}^m$, and conversely adding an element of $q\mathbb{Z}^m$ preserves both membership in Λ and the residue class modulo q .

Therefore, if $\mathbf{w} \leftarrow D_{\Lambda, \sigma}$, then

$$p_{\mathbf{v}} = \Pr[\mathbf{w} \bmod q = \mathbf{v}] = \frac{\rho_{\sigma}(\mathbf{x}_{\mathbf{v}} + q\mathbb{Z}^m)}{\rho_{\sigma}(\Lambda)}. \quad (13)$$

By Lemma 1, applied to the lattice Λ , we have

$$\rho_{\sigma}(\Lambda) = \frac{\sigma^m}{\det(\Lambda)} \rho_{1/\sigma}(\hat{\Lambda}) = \frac{\sigma^m}{q^{n_{\text{dual}}}} \rho_{1/\sigma}(\hat{\Lambda}).$$

On the other hand, applying the Poisson summation formula to the shifted lattice $\mathbf{x}_\mathbf{v} + q\mathbb{Z}^m$ gives

$$\begin{aligned}\rho_\sigma(\mathbf{x}_\mathbf{v} + q\mathbb{Z}^m) &= \frac{\sigma^m}{q^m} \sum_{\mathbf{k} \in \mathbb{Z}^m} \exp\left(-\pi \frac{\sigma^2}{q^2} \|\mathbf{k}\|^2\right) e^{2\pi i \langle \mathbf{k}, \mathbf{x}_\mathbf{v} \rangle / q} \\ &= \frac{\sigma^m}{q^m} \sum_{\mathbf{k} \in \mathbb{Z}^m} \exp\left(-\pi \frac{\sigma^2}{q^2} \|\mathbf{k}\|^2\right) e^{2\pi i \langle \mathbf{k}, \mathbf{v} \rangle / q},\end{aligned}$$

because $\mathbf{x}_\mathbf{v} \equiv \mathbf{v} \pmod{q}$. Combining the two formulas and using $|V| = q^{m-n_{\text{dual}}}$, we obtain

$$p_\mathbf{v} = \frac{1}{|V|} \frac{\sum_{\mathbf{k} \in \mathbb{Z}^m} \exp\left(-\pi \frac{\sigma^2}{q^2} \|\mathbf{k}\|^2\right) e^{2\pi i \langle \mathbf{k}, \mathbf{v} \rangle / q}}{\rho_{1/\sigma}(\hat{A})}. \quad (14)$$

Now set

$$a = \pi \sigma^2 / q^2 = \pi s^2.$$

Since $0 \in \hat{A}$, we have $\rho_{1/\sigma}(\hat{A}) \geq 1$. Also, the phase factor in (14) has absolute value one. Hence the triangle inequality gives

$$p_\mathbf{v} \leq \frac{1}{|V|} \sum_{\mathbf{k} \in \mathbb{Z}^m} e^{-a \|\mathbf{k}\|^2} = \frac{1}{|V|} \left(\sum_{t \in \mathbb{Z}} e^{-at^2} \right)^m.$$

For $t \geq 1$, we have $t^2 \geq t$, and therefore $e^{-at^2} \leq e^{-at}$. Thus

$$\sum_{t \in \mathbb{Z}} e^{-at^2} = 1 + 2 \sum_{t \geq 1} e^{-at^2} \leq 1 + 2 \sum_{t \geq 1} e^{-at} = 1 + \frac{2e^{-a}}{1 - e^{-a}} = \frac{1 + e^{-a}}{1 - e^{-a}}.$$

It follows that

$$p_\mathbf{v} \leq \frac{1}{|V|} \left(\frac{1 + e^{-a}}{1 - e^{-a}} \right)^m = \frac{B}{|V|}.$$

It remains to pass to the distribution conditioned on the event $\mathbf{w} \bmod q \neq \mathbf{0}$. For $\mathbf{v} \in V^\times$,

$$p_\mathbf{v}^* = \frac{p_\mathbf{v}}{1 - p_\mathbf{0}}.$$

The bound above also applies to the zero residue class, so $p_\mathbf{0} \leq B/|V|$. If $B < |V|$, then

$$1 - p_\mathbf{0} \geq 1 - \frac{B}{|V|} = \frac{|V| - B}{|V|}.$$

Therefore, for every $\mathbf{v} \in V^\times$,

$$p_\mathbf{v}^* \leq \frac{B/|V|}{(|V| - B)/|V|} = \frac{B}{|V| - B}.$$

Taking the maximum over $\mathbf{v} \in V^\times$ gives

$$p_{\max}^* = \max_{\mathbf{v} \in V^\times} p_\mathbf{v}^* \leq \frac{B}{|V| - B}.$$

This proves the lemma.

G Proof of Theorem 7

Proof. We decompose the variance into diagonal and off-diagonal terms.

Since $F_{\mathbf{W}} = \sum_{j=1}^N X_j$, after conditioning on $(\mathbf{A}_{\text{dual}}, \mathbf{e})$, we have

$$\begin{aligned} & \text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} (F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})) \\ &= \text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} \left(\sum_{j=1}^N X_j \middle| (\mathbf{A}_{\text{dual}}, \mathbf{e}) \right) \\ &= \sum_{j=1}^N \text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}_j} (X_j | (\mathbf{A}_{\text{dual}}, \mathbf{e})) \\ &\quad + \sum_{j \neq k} \text{Cov}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}_j, \mathbf{W}_k} (X_j, X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})). \end{aligned}$$

Here the second sum is over ordered pairs (j, k) with $j \neq k$.

For the diagonal terms, we only use the boundedness of the individual score. Since $X_j \in [-1, 1]$, we have

$$\begin{aligned} \text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}_j} (X_j | (\mathbf{A}_{\text{dual}}, \mathbf{e})) &\leq \mathbb{E}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}_j} [X_j^2 | (\mathbf{A}_{\text{dual}}, \mathbf{e})] \\ &\leq 1. \end{aligned}$$

Therefore

$$\sum_{j=1}^N \text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}_j} (X_j | (\mathbf{A}_{\text{dual}}, \mathbf{e})) \leq N. \quad (15)$$

For the off-diagonal terms, Theorem 6 gives, for every $j \neq k$,

$$\left| \text{Cov}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}_j, \mathbf{W}_k} (X_j, X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})) - \left(\frac{\kappa_{\text{sq}} + \kappa_{\text{abs}}}{2} - D_{\text{fft}}^2 \right) (h^*(\mathbf{e}))^2 \right| \leq 2(q-1)p_{\text{max}}^*.$$

By the definition of Γ , the triangle inequality therefore implies

$$|\text{Cov}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}_j, \mathbf{W}_k} (X_j, X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e}))| \leq \Gamma.$$

Hence

$$\begin{aligned} & \sum_{j \neq k} \text{Cov}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}_j, \mathbf{W}_k} (X_j, X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e})) \\ &\leq \sum_{j \neq k} |\text{Cov}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}_j, \mathbf{W}_k} (X_j, X_k | (\mathbf{A}_{\text{dual}}, \mathbf{e}))| \leq N(N-1)\Gamma. \end{aligned}$$

Combining this off-diagonal bound with (15) gives

$$\text{Var}_{\mathbf{A}_{\text{fft}}, \mathbf{s}_{\text{fft}}, \mathbf{W}} (F_{\mathbf{W}} | (\mathbf{A}_{\text{dual}}, \mathbf{e})) \leq N + N(N-1)\Gamma.$$

This proves the theorem.

H Proof of Lemma 13

Proof. It suffices to show that the lower bound is increasing as a function of $x = (h^*(\mathbf{e}))^2$.

Denote

$$a := 16 + 16(N-1)\rho, \quad b := 16(N-1)\lambda, \quad c := ND_{\text{ft}}^2.$$

Then $a > 0$, $b \geq 0$, and $c > 0$. Moreover,

$$16 + 16(N-1)\Gamma(x) = a + bx,$$

and

$$16 + 16(N-1)\Gamma(x) + ND_{\text{ft}}^2x = a + (b+c)x.$$

Thus the fractional term in the failure probability is

$$\frac{a + bx}{a + (b+c)x}.$$

For every $x \geq 0$, the denominator is positive. Direct differentiation gives

$$\frac{d}{dx} \left(\frac{a + bx}{a + (b+c)x} \right) = \frac{b(a + (b+c)x) - (b+c)(a + bx)}{(a + (b+c)x)^2} = -\frac{ac}{(a + (b+c)x)^2} \leq 0.$$

Hence this fractional term is decreasing in x , and therefore its negative is increasing in x .

The exponential contribution has the same monotonicity, since

$$\frac{d}{dx} \left(-\exp \left(-\frac{cx}{32} \right) \right) = \frac{c}{32} \exp \left(-\frac{cx}{32} \right) \geq 0.$$

Therefore

$$p_{\text{cor}}(x) = 1 - \frac{a + bx}{a + (b+c)x} - \exp \left(-\frac{cx}{32} \right)$$

is increasing for $x \geq 0$.

Finally, on the norm good event, we have

$$h^*(\mathbf{e}) \geq h^*(u_*).$$

Since the lower bound $h^*(u_*)$ is nonnegative in the parameter regime under consideration, this implies

$$(h^*(\mathbf{e}))^2 \geq (h^*(u_*))^2.$$

The claimed lower bound follows by applying the monotonicity of $p_{\text{cor}}(x)$.

I Proof of Lemma 14

Proof. The incorrect guess creates a nonzero linear phase in the fresh random block $\mathbf{A}_{\text{guess}}$. All remaining terms are collected into the auxiliary random variable C , which is independent of $\mathbf{A}_{\text{guess}}$. By the law of total expectation, it is enough to prove that, for every fixed real number c ,

$$\mathbb{E}_{\mathbf{A}_{\text{guess}}} \left[\cos \left(\frac{2\pi}{q} (\Delta^T \mathbf{A}_{\text{guess}} \delta + c) \right) \right] = 0.$$

We first show that the vector $\Delta^T \mathbf{A}_{\text{guess}}$ is uniform over $\mathbb{Z}_q^{n_{\text{guess}}}$. Let $\mathbf{a}_j$ be the j -th column of $\mathbf{A}_{\text{guess}}$. Since $\Delta \neq \mathbf{0}$, there exists an index ℓ such that $\Delta_\ell \neq 0$. If all coordinates of $\mathbf{a}_j$ except the ℓ -th coordinate are fixed, then

$$\Delta^T \mathbf{a}_j = \Delta_\ell a_{\ell,j} + \text{constant}.$$

Because q is prime and $\Delta_\ell \neq 0$, multiplication by Δ_ℓ is a bijection on $\mathbb{Z}_q$. Thus, as $a_{\ell,j}$ ranges uniformly over $\mathbb{Z}_q$, the value $\Delta^T \mathbf{a}_j$ is also uniform over $\mathbb{Z}_q$. This argument applies to each column separately. Since the columns of $\mathbf{A}_{\text{guess}}$ are independent, the random vector

$$\mathbf{y}^T := \Delta^T \mathbf{A}_{\text{guess}}$$

is uniform over $\mathbb{Z}_q^{n_{\text{guess}}}$.

We next use the fact that the wrong guess gives a nonzero difference vector δ . Since $\delta \neq \mathbf{0}$, choose an index j_0 such that $\delta_{j_0} \neq 0$. For a uniform vector $\mathbf{y} \in \mathbb{Z}_q^{n_{\text{guess}}}$, if all coordinates except y_{j_0} are fixed, then

$$\mathbf{y}^T \delta = \delta_{j_0} y_{j_0} + \text{constant}.$$

Again, multiplication by the nonzero element δ_{j_0} is a bijection on $\mathbb{Z}_q$. Hence $\mathbf{y}^T \delta$ is uniform over $\mathbb{Z}_q$. Applying this to $\mathbf{y}^T = \Delta^T \mathbf{A}_{\text{guess}}$, we obtain

$$\Delta^T \mathbf{A}_{\text{guess}} \delta = \mathbf{y}^T \delta$$

uniformly distributed over $\mathbb{Z}_q$.

Therefore, for fixed c , the cosine average is

$$\begin{aligned} & \mathbb{E}_{\mathbf{A}_{\text{guess}}} \left[\cos \left(\frac{2\pi}{q} (\Delta^T \mathbf{A}_{\text{guess}} \delta + c) \right) \right] \\ &= \frac{1}{q} \sum_{u \in \mathbb{Z}_q} \cos \left(\frac{2\pi}{q} (u + c) \right) \\ &= \Re \left(e^{2\pi i c / q} \cdot \frac{1}{q} \sum_{u \in \mathbb{Z}_q} e^{2\pi i u / q} \right) = 0, \end{aligned}$$

because the sum of all nontrivial q -th roots of unity is zero. Finally, taking expectation over the auxiliary random variable C gives

$$\mathbb{E}_{\mathbf{A}_{\text{guess}}, C} \left[\cos \left(\frac{2\pi}{q} (\Delta^T \mathbf{A}_{\text{guess}} \delta + C) \right) \right] = 0.$$

This proves the claim.