

Provable dual attack on LWE via lattice projection

Hongyuan Qu^{1,2}, Chengliang Tian³, Geng Wang⁴, and Guangwu Xu^{1,2(✉)}

¹ Key Laboratory of Cryptologic Technology and Information Security of
Ministry of Education, Qingdao, Shandong, China
`qhy1qhy@mail.sdu.edu.cn, gxu4sdq@sdu.edu.cn`

² School of Cyber Science and Technology, Shandong University,
Qingdao, Shandong, China.

³ College of Computer Science and Technology, Qingdao University, China
`tianchengliang@qdu.edu.cn`

⁴ School of Computer Science, Shanghai Jiao Tong University,
Shanghai, China
`wanggxx@sjtu.edu.cn`

Abstract. The Learning with Errors (LWE) problem is a cornerstone of post-quantum cryptography, and the dual attack is a central tool for evaluating its concrete hardness. In provable dual attacks, a major bottleneck is the large number of discrete Gaussian samples required over the dual lattice. While recent works have focused on accelerating the guessing step via modulus switching, the sampling bottleneck has remained largely unoptimized. In this paper, we address this issue from a new perspective. First, we observe that Banaszczyk’s inequality can be significantly improved when the first minimum of the lattice is suitably large. This yields a refined distinguishing inequality, which provides sharper conditions for attack success. Second, we propose a new provable dual attack framework based on lattice projection. Instead of sampling over the entire dual lattice, we sample over a carefully chosen sublattice, whose dual is the projection of the primal lattice onto a lower-dimensional subspace. This reduces the lattice dimension and thus substantially accelerates discrete Gaussian sampling. We rigorously analyze the parameter conditions for the success of our framework, and by enumerating all feasible parameters we achieve an improvement of 26.1–43.7 bits over the original provable dual attack framework.

1 Introduction

Since its introduction by Regev in 2005, the Learning with Errors (LWE) problem [25] has served as a fundamental mathematical hard problem in modern cryptography. Owing to its resistance to quantum attacks, scalability, and flexibility, it has become a cornerstone for constructing post-quantum secure cryptographic protocols. LWE underpins not only advanced cryptographic primitives such as digital signatures [10, 12, 31], identity-based encryption (IBE) [14], non-interactive zero-knowledge proofs (NIZK) [22, 17], and fully homomorphic

encryption (FHE) [5, 7, 8], but also holds a central position in the NIST Post-Quantum Cryptography (PQC) standardization process. For example, the selected algorithms Kyber [4] and Dilithium [10] both rely on algebraic variants of LWE for their security, and their parameter choices depend directly on the security evaluation of LWE. Thus, concrete hardness analysis of LWE acts as a critical bridge between theoretical post-quantum cryptography and engineering implementations.

This paper focuses on the Search-LWE problem: given a random matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$, where the error vector $\mathbf{e} \in \mathbb{Z}_q^m$ is drawn from a discrete Gaussian distribution, the objective is to recover the secret $\mathbf{s} \in \mathbb{Z}_q^n$.

A key theoretical result by Regev [25] shows that the LWE problem is provably hard under quantum reductions from worst-case lattice problems, such as the Shortest Vector Problem (SVP) and Bounded Distance Decoding (BDD). This worst-case hardness provides strong security guarantees for constructions based on LWE. Nonetheless, concrete parameter selection—particularly the dimension n , modulus q , and noise width σ —critically influences practical security and efficiency.

The deployment of LWE in practical cryptosystems has led to a co-evolution of parameter sizes and cryptanalytic techniques. Early cryptanalysis primarily relied on primal attacks, which reduce LWE to the Unique Shortest Vector Problem (uSVP) and apply lattice reduction algorithms like BKZ. However, the super-polynomial complexity of BKZ in high dimensions limits its practicality, motivating the development of dual attacks. These attacks distinguish LWE samples from random by examining short vectors in the dual lattice, showing particular promise in high-dimensional regimes.

Recent advances have significantly improved the efficiency of dual attacks, in some cases surpassing primal approaches. Progress has been made in two main directions: (1) improved short-vector generation using lattice sieving [13, 20, 3] or re-randomization from BKZ-reduced bases [15], and (2) accelerated guessing via the Discrete Fourier Transform (DFT) to compute scores simultaneously [15]. Further efficiency gains were achieved by incorporating modulus switching—a technique from fully homomorphic encryption—to reduce DFT costs [18]. Different modulus switching strategies were adopted in [15] and [18].

Despite these improvements, such attacks often depend on heuristic assumptions, such as short-vector uniformity and noise inner-product independence. Ducas et al. [11] showed that these heuristics can lead to attack failure in certain settings, casting doubt on their reliability. Subsequent work by Carrier et al. [6] replaced modulus switching with linear code decoding, thereby avoiding some heuristic assumptions but still relying on unproven ones.

A major step toward rigorous analysis was taken by [23], who proposed the first heuristic-free dual attack framework with provable guarantees. Their attack effectively avoids the paradoxical regimes identified in [11]. Nevertheless, integrating modulus switching into this provable framework remained an open problem. [24] resolved this open problem by employing rational q -ary lattices to eliminate the additional error introduced during the modulus switching process,

thereby proposing a provable dual attack method that incorporates modulus switching; they also refined the distinguishing inequality used in [23]. More recently, [29] proposed a p -adic layered modulus switching approach for provable dual attacks. Instead of recovering residues modulo several distinct primes, their method fixes a single small prime and recovers the guessed secret digit by digit in its p -adic expansion, which reduces the dominant guessing term from the largest CRT prime to that of a fixed small prime.

We observe that both [24] and [29] focus on the modulus switching procedure, which accelerates the guessing step of the attack. However, the primary bottleneck of provable dual attacks—discrete Gaussian sampling—has not been directly optimized. In this paper, we exploit the idea of projection: by projecting the original dual attack into a subspace, we reduce the lattice dimension and thereby significantly improve the efficiency of discrete Gaussian sampling.

1.1 Our Contributions

Our contributions are two-fold. First, we observe that Banaszczyk’s inequality can be further significantly improved in certain cases. Specifically, by imposing the additional condition that $\lambda_1(L)$ is suitably large, we obtain an improved inequality (Lemma 7). This yields a refined distinguishing inequality (Corollary 3): for typical parameters with $n \geq 500$, if $r_{\mathbf{x}} \geq cs\sqrt{n/(2\pi)}$ and $\lambda_1(L) \geq kcs\sqrt{n/(2\pi)}$ with $c \geq 1$ and $k \geq 1.04$, then

$$\rho_s(r_{\mathbf{x}}) \leq \frac{\rho_s(L + \mathbf{x})}{\rho_s(L)} \leq 2 \left(e^{1-c^2} \right)^{n/2}.$$

Compared with Banaszczyk’s inequality, this gives an improvement by an exponential factor for $c > 1$. This refined inequality provides sharper conditions for attack success and serves as a key ingredient in our provable dual attack framework.

Our second contribution is a new provable dual attack framework via the idea of lattice projection. Specifically, a provable dual attack essentially amounts to distinguishing whether the distance from a point $\mathbf{t}$ to a lattice L is small or large. This requires many discrete Gaussian samples over the dual lattice L^* , which is the main bottleneck of this approach. Previous works [23, 24, 29] sample over the entire L^* . In contrast, we adopt the following simple idea: we sample over a sublattice of L^* . More precisely, let $\mathbf{c}_1, \dots, \mathbf{c}_n$ be a BKZ-reduced basis of L^* , choose a suitable $m < n$, and let L' be the sublattice generated by $\mathbf{c}_1, \dots, \mathbf{c}_m$. We perform discrete Gaussian sampling over L' . Since m is typically much smaller than n , this significantly accelerates discrete Gaussian sampling. A key observation is that the dual of L' is $\pi(L)$, where $\pi : \mathbb{R}^n \rightarrow \text{span}(L')$ is the orthogonal projection. Consequently, the original problem of distinguishing the distance from $\mathbf{t}$ to L is transformed into distinguishing the distance from $\pi(\mathbf{t})$ to $\pi(L)$. Moreover, using the property that $\langle \mathbf{t}, \mathbf{v} \rangle = \langle \pi(\mathbf{t}), \mathbf{v} \rangle$ for all $\mathbf{v} \in \text{span}(L')$, we can still perform FFT over $\mathbb{Z}_q$ to accelerate the attack. With the help of our improved Banaszczyk’s inequality, we rigorously analyze the parameter conditions for the

success of our attack framework. Despite the unavoidable use of the most common assumptions in lattice cryptanalysis, namely the Gaussian Heuristic (GH) and the Geometric Series Assumption (GSA), we enumerate all feasible parameters and take the optimal values, achieving an improvement of 26.1–43.7 bits over the original framework [23].

1.2 Roadmap

Section 2 introduces some necessary background knowledge. Section 3 presents the improved Banaszczyk’s inequality we obtained. Section 4 introduces our provable dual attack framework, together with its parameter selection and experimental results. Section 5 concludes the paper and points out possible directions for future research.

2 Preliminaries

Let $\mathbb{N}, \mathbb{Z}, \mathbb{R}$ and $\mathbb{C}$ denote the sets of natural numbers, integers, real numbers, and complex numbers, respectively. We use bold uppercase letters (e.g., $\mathbf{A}$) for matrices and bold lowercase letters (e.g., $\mathbf{v}$) for column vectors. For a vector $\mathbf{v} \in \mathbb{R}^n$, we denote by v_i its i -th component for $i = 1, \dots, n$. For two vectors $\mathbf{a}$ and $\mathbf{b}$, their concatenation is denoted as $(\mathbf{a}, \mathbf{b})$. Unless stated otherwise, $\log$ refers to the base-2 logarithm. For a positive integer q , let $\mathbb{Z}_q$ represent the additive group of integers modulo q . Elements of $\mathbb{Z}_q$ are represented by their least non-negative residues $\{0, 1, \dots, q-1\}$. For a complex number z , $\text{Re}(z)$ and $|z|$ denote its real part and modulus, respectively. For a matrix $\mathbf{A}$, $\mathbf{A}^T$ is its transpose and $\text{rk}(\mathbf{A})$ denotes its rank. The cardinality of a finite set X is written as $|X|$. Given sets $X_1, \dots, X_n$, their union is denoted by $\bigcup_{i=1}^n X_i$. If they are pairwise disjoint, their disjoint union is written as $\bigsqcup_{i=1}^n X_i$.

For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, their inner product is written as $\langle \mathbf{x}, \mathbf{y} \rangle$, and the Euclidean norm of $\mathbf{x}$ is $\|\mathbf{x}\|$. The open ball of radius $r > 0$ centered at $\mathbf{x}$ is defined as $\mathcal{B}(\mathbf{x}, r) = \{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{y} - \mathbf{x}\| < r\}$. We use $\mathcal{B}(r)$ to represent $\mathcal{B}(\mathbf{0}, r)$ when $\mathbf{x} = \mathbf{0}$.

The uniform distribution over a set X is denoted by $\mathcal{U}(X)$. We write $x \leftarrow \chi$ to indicate that x is sampled from a distribution χ over a domain which is clear from context. For two distributions P and Q on a countable set X , their total variation distance is $d_{TV}(P, Q) = \frac{1}{2} \sum_{x \in X} |P(x) - Q(x)|$. This metric satisfies the triangle inequality: $d_{TV}(P, S) \leq d_{TV}(P, Q) + d_{TV}(Q, S)$ for any distributions P, Q, S . It also satisfies the data-processing inequality: for any function f , $d_{TV}(f(P), f(Q)) \leq d_{TV}(P, Q)$. We will use the following standard concentration inequality.

Theorem 1 (Hoeffding’s inequality). *Let $X_1, \dots, X_N$ be independent random variables such that $a_i \leq X_i \leq b_i$. Consider the sum $S_N = X_1 + \dots + X_N$. Then*

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

for all $t > 0$.

2.1 Lattices

A lattice in $\mathbb{R}^n$ is a discrete subgroup of the additive group $(\mathbb{R}^n, +)$. Here, discreteness means that there exists a positive d such that $\mathcal{B}(\mathbf{0}, d) \cap L = \{\mathbf{0}\}$. Every lattice $L \subset \mathbb{R}^n$ has a basis: there are linearly independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_r \in L$ such that

$$L = \{u_1 \mathbf{b}_1 + \dots + u_r \mathbf{b}_r : u_1, \dots, u_r \in \mathbb{Z}\}.$$

The integer r is the rank of L and n is its dimension. L is called full-rank when $r = n$. If $\mathbf{B}$ is a basis matrix of L , the determinant of L is $\det(L) = \sqrt{\det(\mathbf{B}^T \mathbf{B})}$. The length of a shortest nonzero vector in L is denoted by $\lambda_1(L) = \min\{\|\mathbf{v}\| : \mathbf{v} \in L \setminus \{\mathbf{0}\}\}$. In practice, $\lambda_1(L)$ is often approximated by the Gaussian heuristic $GH(L) = \sqrt{\frac{n}{2\pi e}} \det(L)^{1/n}$. The dual lattice of L is

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

where $\text{span}(L)$ denotes the linear span of L . One has $(L^*)^* = L$. If L is a full-rank lattice and $\mathbf{B}$ is a basis of L , then a basis of L^* is $(\mathbf{B}^{-1})^T$. For a point $\mathbf{x} \in \mathbb{R}^n$, its distance to lattice L is

$$\text{dist}(\mathbf{x}, L) = \min\{\|\mathbf{v} - \mathbf{x}\| : \mathbf{v} \in L\}.$$

The following lemma plays a key role in our attack algorithm.

Lemma 1. *Let $L \subset \mathbb{R}^n$ be a full-rank lattice. Let $\mathbf{c}_1, \dots, \mathbf{c}_n$ be a lattice basis of L^* . Let $m < n$, and let $L' = L(\mathbf{c}_1, \dots, \mathbf{c}_m)$ be the sublattice with lattice basis the first m vectors. Set $V = \text{span}(L')$. Then*

$$(L')^* = \pi(L),$$

where $\pi : \mathbb{R}^n \rightarrow V$ is the orthogonal projection of $\mathbb{R}^n$ onto V .

Proof. Since $L' \subseteq L^*$, for every $\mathbf{v} \in L$ and every $\mathbf{u} \in L'$, we have $\langle \mathbf{v}, \mathbf{u} \rangle \in \mathbb{Z}$. Hence

$$\langle \pi(\mathbf{v}), \mathbf{u} \rangle = \langle \mathbf{v}, \mathbf{u} \rangle \in \mathbb{Z},$$

and $\pi(\mathbf{v}) \in V$. Thus $\pi(L) \subseteq (L')^*$, so $\pi(L) \subseteq (L')^*$.

Conversely, let $\mathbf{b}_1, \dots, \mathbf{b}_n$ be a lattice basis of L satisfying

$$\langle \mathbf{b}_i, \mathbf{c}_j \rangle = \delta_{i,j}$$

for all $i, j \in \{1, \dots, n\}$. For any $\mathbf{t} \in (L')^*$, let $V^\perp$ be the orthogonal complement of V in $\mathbb{R}^n$. Then $\mathbf{b}_{m+1}, \dots, \mathbf{b}_n$ form a basis of $V^\perp$. Write

$$\mathbf{t} = \sum_{i=1}^n \alpha_i \mathbf{b}_i, \quad \alpha_i \in \mathbb{R}.$$

Then for every $i \in \{1, \dots, m\}$,

$$\alpha_i = \langle \mathbf{t}, \mathbf{c}_i \rangle \in \mathbb{Z}.$$

Therefore

$$\mathbf{t} - \sum_{i=m+1}^n \alpha_i \mathbf{b}_i \in L.$$

Consequently,

$$\pi(\mathbf{t}) = \pi \left(\mathbf{t} - \sum_{i=m+1}^n \alpha_i \mathbf{b}_i \right) \in \pi(L),$$

since $\sum_{i=m+1}^n \alpha_i \mathbf{b}_i \in V^\perp$. Since $\mathbf{t} \in V$, i.e. $\mathbf{t} = \pi(\mathbf{t})$, it follows that $\mathbf{t} \in \pi(L)$. Therefore $(L')^* = \pi(L)$. $\square$

2.2 LWE

Introduced by Regev [25], the Learning With Errors (LWE) problem is a computationally hard lattice problem that serves as the security foundation for numerous modern cryptographic schemes. Formally, the associated distribution is defined as follows.

Definition 1 (LWE distribution). *Let $n, m, q \in \mathbb{N}$, let q be a prime, 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 entries of the matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ independently and uniformly over $\mathbb{Z}_q$, sampling $\mathbf{s} \in \mathbb{Z}_q^n$ and $\mathbf{e} \in \mathbb{Z}_q^m$ independently from χ_s and χ_e , respectively, and outputting $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q)$.*

In this paper, we mainly consider the search version of the LWE problem, that is, given a sample from $\text{LWE}_{n,m,q,\chi_s,\chi_e}$, recover $\mathbf{s}$.

In recent years, when designing LWE-based cryptographic schemes, χ_s and χ_e have often been chosen as small distributions. For example, in the NIST post-quantum standardization scheme Kyber, both χ_s and χ_e adopt the following centered binomial distribution:

Definition 2 (Centered binomial distribution). *Let $\alpha \in \{0, \dots, \frac{q-1}{2}\}$. The centered binomial distribution B_α is defined by*

$$B_\alpha \sim \sum_{i=1}^{\alpha} (X_i - Y_i),$$

where X_i and Y_i are i.i.d. uniform random variables over $\{0, 1\}$. In particular, for all $i \in \{-\alpha, \dots, \alpha\}$,

$$\Pr(B_\alpha = i) = 2^{-2\alpha} \binom{2\alpha}{\alpha + i}.$$

B_α has mean 0 and standard deviation $\sigma = \sqrt{\frac{\alpha}{2}}$.

Note that Kyber uses three parameter sets, denoted by Kyber512, Kyber768, and Kyber1024. In Kyber512, the centered binomial distribution parameter is $\alpha = 3$. In Kyber768 and Kyber1024, the parameter is $\alpha = 2$.

2.3 Fourier Transform

The Fourier transform of a rapidly decreasing smooth function⁵ $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined by

$$\forall \mathbf{w} \in \mathbb{R}^n, \hat{h}(\mathbf{w}) = \int_{\mathbb{R}^n} h(\mathbf{x}) \exp(-2\pi i \langle \mathbf{x}, \mathbf{w} \rangle) d\mathbf{x}.$$

Let G be a finite abelian group and $\hat{G}$ be its dual group. For a function $f : G \rightarrow \mathbb{C}$, its Fourier transform $\hat{f} : \hat{G} \rightarrow \mathbb{C}$ is defined by

$$\hat{f}(\chi) := \sum_{g \in G} f(g) \chi(g),$$

where $\chi \in \hat{G}$. Specifically, for a function $f : \mathbb{Z}_q^n \rightarrow \mathbb{C}$, its Fourier transform $\hat{f}$ is given by

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

The Fourier transform of a function $f : \mathbb{Z}_q^n \rightarrow \mathbb{C}$ can be calculated in time $O(n \log q \cdot q^n)$ using the Fast Fourier Transform (FFT) algorithm. We need the following concrete FFT complexity:

Lemma 2 ([6, Proposition 2.11]). *Let $q = 3329$ and n be a positive integer. There exists an FFT over $\mathbb{Z}_q^n$ with complexity given by*

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

where

$$C_{\text{add}} = 160, \quad C_{\text{mul}} = 1024, \quad N_{\text{FFT}}^{(\text{add})} = 240500, \quad N_{\text{FFT}}^{(\text{mul})} = 115928.$$

2.4 Discrete Gaussian Measure over Lattices

For a parameter $s > 0$ and a vector $\mathbf{x} \in \mathbb{R}^n$, the Gaussian function is defined as $\rho_s(\mathbf{x}) = e^{-\frac{\pi \|\mathbf{x}\|^2}{s^2}}$. Given a lattice $L \subset \mathbb{R}^n$, we set $\rho_s(L) = \sum_{\mathbf{v} \in L} \rho_s(\mathbf{v})$. The discrete Gaussian distribution on L with parameter s , denoted by $D_{L,s}$, is defined so that for every $\mathbf{v} \in L$,

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

We further introduce the function $f_{L,s} : \mathbb{R}^n \rightarrow \mathbb{R}$ given by

$$f_{L,s}(\mathbf{t}) = \frac{\rho_s(L + \mathbf{t})}{\rho_s(L)}, \tag{1}$$

which always satisfies $0 < f_{L,s}(\mathbf{t}) \leq 1$.

A key tool is Poisson summation formula, which we state as follows:

⁵ This means that h and all its (partial) derivatives $D^\beta h$ are rapidly decreasing in the sense that $\sup_{\mathbf{x} \in \mathbb{R}^n} |\mathbf{x}^\alpha D^\beta h(\mathbf{x})| < \infty$ for every $\alpha, \beta \in \mathbb{N}^n$. Such a function is said to be in the Schwartz space [27].

Lemma 3 (Poisson summation formula[26, Chapter VII, Proposition 15]). *For an n -dimensional lattice L , let $s > 0$ and $\mathbf{t} \in \text{span}(L)$, the following hold:*

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

Using Poisson summation formula, the function $f_{L,s}$ can be rewritten as an expectation over the dual lattice:

$$\begin{aligned} f_{L,s}(\mathbf{t}) &= \frac{\rho_s(L + \mathbf{t})}{\rho_s(L)} = \frac{\sum_{\mathbf{w} \in L^*} \exp(2\pi i \langle \mathbf{t}, \mathbf{w} \rangle) \rho_{1/s}(\mathbf{w})}{\rho_{1/s}(L^*)} \\ &= \mathbb{E}_{\mathbf{w} \leftarrow D_{L^*, 1/s}}[\exp(2\pi i \langle \mathbf{t}, \mathbf{w} \rangle)] \end{aligned} \quad (2)$$

is in fact the expectation of a random variable defined on L^* .

Next, we introduce the classic measure inequality of Banaszczyk [2, Lemma 1.5]. The following statement is taken from [28] (Theorem 3.1 and its remark) which is an improvement of the Banaszczyk inequality with the factor 2 removed.

Lemma 4. *Let $L \subset \mathbb{R}^n$ be a lattice of dimension n . Then for any $c \geq 1$ and $\mathbf{t} \in \mathbb{R}^n$,*

$$\frac{\rho_s(L + \mathbf{t}) \setminus \mathcal{B}(\mathbf{0}, cs\sqrt{\frac{n}{2\pi}})}{\rho_s(L)} \leq \left(c\sqrt{e} \cdot e^{-\frac{c^2}{2}}\right)^n.$$

Let $\tau = s\sqrt{\frac{n}{2\pi}}$. For an n -dimensional lattice L and a vector $\mathbf{x} \in \mathbb{R}^n$ let $r_{\mathbf{x}} = \text{dist}(\mathbf{x}, L)$. The next corollary provides upper and lower bounds for $f_{L,s}(\mathbf{x})$, which plays a crucial role in the formal analysis and parameter selection of provable dual attack in [24].

Corollary 1 (Distinguishing inequality[24, Corollary 1]). *Let $L \subset \mathbb{R}^n$ be an n -dimensional lattice and $\mathbf{x} \in \mathbb{R}^n$. If $r_{\mathbf{x}} \geq \tau$, then*

$$\rho_s(r_{\mathbf{x}}) \leq f_{L,s}(\mathbf{x}) \leq \left(\frac{r_{\mathbf{x}}}{\tau}\right)^n e^{\frac{n}{2}} \rho_s(r_{\mathbf{x}}).$$

Moreover, the right-hand side is a decreasing function of $r_{\mathbf{x}}$.

The dual attack relies heavily on lattice reduction algorithms such as BKZ to find several short vectors in the dual lattice. BKZ has the following complexity result, originally proposed in [18, Lemma 4.1, Assumption 7.3]:

Lemma 5 (BKZ Complexity [6, Lemma 2.7]). *Let B be a basis of a d -dimensional lattice. Then the running time T_{BKZ} of the BKZ algorithm with block size β is*

$$T_{\text{BKZ}}(n, \beta) = C_{\text{prog}}^2 \cdot (n - \beta + 1) \cdot T_{\text{NNS}}(\beta^{\text{eff}}),$$

where

- $T_{\text{NNS}}(\beta)$ is the time complexity for finding all close pairs in dimension β ,

- $C_{\text{prog}} = 1/(1 - 2^{-0.292})$ is the number of close pairs searches to run,
- β^{eff} is the optimal sieve dimension to use for solving the Shortest Vector Problem (SVP) for lattices in dimension β , and it is estimated in [9] that $\beta^{\text{eff}} = \beta - \frac{\beta \log(4/3)}{\log(\beta/(2\pi e))}$.

Under the GH and GSA assumptions, for any integer $0 < m < n$, we have the following result:

Lemma 6. *For any integer $0 < m < n$, let L_m be the sublattice of L with basis $\mathbf{b}'_1, \dots, \mathbf{b}'_m$. Then*

$$\det(L_m) = \delta(\beta)^{m(n-m)} \det(L)^{\frac{m}{n}}.$$

Given a lattice $L \subset \mathbb{R}^n$ with basis $\mathbf{B}$ and a parameter $s > 0$, the discrete Gaussian sampling algorithm takes $\mathbf{B}$ and s as inputs and outputs a vector sampled from the distribution $D_{L,s}$. The sampling complexity increases as s decreases. The smoothing parameter [19] quantifies this difficulty: polynomial-time sampling algorithms exist when s is larger than the smoothing parameter [16, 14, 21], while sampling becomes significantly more difficult when s falls below the smoothing parameter. Notably, for sufficiently small s , the sampling algorithm can be used to solve SVP [1]. Among existing methods, the Monte Carlo Markov Chain (MCMC)-based algorithm in [30] achieves sampling for all $s > 0$ and will be adopted in this work. The following theorem demonstrates the time complexity of this sampling algorithm.

Theorem 2 ([30, Theorem 1]). *There is an algorithm that, given a basis $\mathbf{B}$ of a rank- m lattice $L \subset \mathbb{R}^n$, any $\epsilon > 0$, and any $s > 0$, returns a sample according to some distribution $D_{L,s,\epsilon}$ such that $d_{TV}(D_{L,s,\epsilon}, D_{L,s}) \leq \epsilon$. This algorithm runs in time $\ln(\frac{1}{\epsilon}) \cdot \frac{1}{\Delta} \cdot \text{poly}(n)$, where*

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

and $\mathbf{b}_1^*, \dots, \mathbf{b}_n^*$ are the Gram-Schmidt vectors of $\mathbf{B}$.

3 Optimization of Banaszczyk's Inequality and Improved Distinguishing Inequality

We observe that Banaszczyk's inequality can be further significantly improved in certain cases. Specifically, we impose an additional condition that $\lambda_1(L)$ is suitably large and obtain the following improved inequality:

Lemma 7 (Improved Banaszczyk's Inequality). *Let $L \subset \mathbb{R}^n$ be a lattice. Assume that $\lambda_1(L) \geq kcs\sqrt{\frac{n}{2\pi}}$ where $k > 1$. Then for any $c \geq 1$ and $\mathbf{t} \in \mathbb{R}^n$,*

$$\rho_s\left((L + \mathbf{t}) \setminus \mathcal{B}\left(\mathbf{0}, cs\sqrt{\frac{n}{2\pi}}\right)\right) \leq \frac{\left(e^{1-c^2}\right)^{\frac{n}{2}}}{1 - \epsilon},$$

where $\epsilon = \left(k\sqrt{e} \cdot e^{-\frac{k^2}{2}}\right)^n$.

Proof. We have

$$\begin{aligned} \rho_s \left((L + \mathbf{t}) \setminus \mathcal{B} \left(\mathbf{0}, cs \sqrt{\frac{n}{2\pi}} \right) \right) &= \sum_{\substack{\mathbf{v} \in L + \mathbf{t} \\ \|\mathbf{v}\| \geq cs \sqrt{\frac{n}{2\pi}}}} e^{-\frac{\pi}{s^2} \left(1 - \frac{1}{c^2}\right) \|\mathbf{v}\|^2} \cdot e^{-\frac{\pi}{s^2 c^2} \|\mathbf{v}\|^2} \\ &\leq e^{-\frac{\pi}{s^2} \left(1 - \frac{1}{c^2}\right) \frac{c^2 s^2 n}{2\pi}} \cdot \rho_{sc}(L + \mathbf{t}) \leq \left(e^{1-c^2} \right)^{\frac{n}{2}} \cdot \rho_{sc}(L). \end{aligned}$$

Since $\lambda_1(L) \geq kcs\sqrt{\frac{n}{2\pi}}$, using Lemma 4 we obtain

$$\frac{\rho_{sc}(L) - 1}{\rho_{sc}(L)} = \frac{\rho_{sc}(L \setminus \mathcal{B}(\mathbf{0}, kcs\sqrt{\frac{n}{2\pi}}))}{\rho_{sc}(L)} \leq \left(k\sqrt{e} \cdot e^{-\frac{k^2}{2}} \right)^n.$$

Let $\epsilon = \left(k\sqrt{e} \cdot e^{-\frac{k^2}{2}} \right)^n$. For $k > 1$, $\epsilon < 1$. Then $\rho_{sc}(L) \leq \frac{1}{1-\epsilon}$. Thus we have

$$\rho_s \left((L + \mathbf{t}) \setminus \mathcal{B} \left(\mathbf{0}, cs \sqrt{\frac{n}{2\pi}} \right) \right) \leq \left(e^{1-c^2} \right)^{\frac{n}{2}} \cdot \rho_{sc}(L) \leq \frac{\left(e^{1-c^2} \right)^{\frac{n}{2}}}{1 - \epsilon}.$$

Since $\rho_s(L) > 1$ for any $s > 0$, the following is immediate from Lemma 7:

Corollary 2. *Let $L \subset \mathbb{R}^n$ and $s > 0$. Let $\mathbf{x} \in \mathbb{R}^n$ and $r_{\mathbf{x}} = \text{dist}(\mathbf{x}, L)$. Assume $r_{\mathbf{x}} \geq cs\sqrt{\frac{n}{2\pi}}$ and $\lambda_1(L) \geq kcs\sqrt{\frac{n}{2\pi}}$ for $c \geq 1$ and $k > 1$. Then*

$$\rho_s(r_{\mathbf{x}}) \leq \frac{\rho_s(L + \mathbf{x})}{\rho_s(L)} \leq \frac{\left(e^{1-c^2} \right)^{\frac{n}{2}}}{1 - \epsilon},$$

where $\epsilon = \left(k\sqrt{e} \cdot e^{-\frac{k^2}{2}} \right)^n$.

In a typical application, $n \geq 500$ (e.g., [18, 6]). The distinguishing inequality can be improved to

Corollary 3 (Improved distinguishing inequality). *Let $L \subset \mathbb{R}^n$ and $s > 0$. Let $\mathbf{x} \in \mathbb{R}^n$ and $r_{\mathbf{x}} = \text{dist}(\mathbf{x}, L)$. Assume $r_{\mathbf{x}} \geq cs\sqrt{\frac{n}{2\pi}}$ and $\lambda_1(L) \geq kcs\sqrt{\frac{n}{2\pi}}$ where $c \geq 1$ and $k \geq 1.04$. Assume $n \geq 500$. Then*

$$\rho_s(r_{\mathbf{x}}) \leq \frac{\rho_s(L + \mathbf{x})}{\rho_s(L)} \leq 2 \left(e^{1-c^2} \right)^{\frac{n}{2}}.$$

Proof. When $k \geq 1.04$ and $n \geq 500$, we have

$$\left(k\sqrt{e} \cdot e^{-\frac{k^2}{2}} \right)^n < 0.5.$$

Then the result follows from Corollary 2.

Remark 1. Note that in Corollary 3, the term $2 \left(e^{1-c^2} \right)^{\frac{n}{2}}$ can equivalently be expressed as $2e^{\frac{n}{2}} e^{\frac{\pi}{s^2}(r_{\mathbf{x}}^2 - r^2 c^2)} \rho_s(r_{\mathbf{x}})$. Hence, the above improved distinguishing inequality can be written equivalently in the following unified form, where the same factor $\rho_s(r_{\mathbf{x}})$ appears on both sides:

$$\rho_s(r_{\mathbf{x}}) \leq \frac{\rho_s(L + \mathbf{x})}{\rho_s(L)} \leq 2e^{\frac{n}{2}} e^{\frac{\pi}{s^2}(r_{\mathbf{x}}^2 - r^2 c^2)} \rho_s(r_{\mathbf{x}}).$$

Observe that

$$\frac{\left(c\sqrt{e} \cdot e^{-\frac{c^2}{2}} \right)^n}{2 \left(e^{1-c^2} \right)^{\frac{n}{2}}} = \frac{c^n}{2}.$$

Here, the numerator is the upper bound given by Banaszczyk's inequality. For $c > 1$, we obtain an improvement of Banaszczyk's inequality by an exponential factor.

4 Provable Dual Attack via Projection onto the Span of a Dual Sublattice

4.1 Overview of our Algorithm

Given $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$, where $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$, $m \geq n$, and $\mathbf{A}$ has full rank over $\mathbb{Z}_q$, we partition the coordinates $I := \{1, \dots, n\}$ of $\mathbf{s}$ into three parts I_{enu} , I_{fft} , and I_{lat} , of sizes n_{enu} , n_{fft} , and n_{lat} , respectively. Correspondingly, $\mathbf{s}$ is divided into three parts $\mathbf{s}_{\text{enu}}$, $\mathbf{s}_{\text{fft}}$, and $\mathbf{s}_{\text{lat}}$. Similarly, $\mathbf{A}$ is divided into three parts $\mathbf{A}_{\text{enu}}$, $\mathbf{A}_{\text{fft}}$, and $\mathbf{A}_{\text{lat}}$. Then

$$\mathbf{b} = \mathbf{A}_{\text{enu}}\mathbf{s}_{\text{enu}} + \mathbf{A}_{\text{fft}}\mathbf{s}_{\text{fft}} + \mathbf{A}_{\text{lat}}\mathbf{s}_{\text{lat}} + \mathbf{e} \bmod q.$$

Thus

$$\begin{bmatrix} \mathbf{b} - \mathbf{A}_{\text{fft}}\mathbf{s}_{\text{fft}} \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{\text{lat}} & q\mathbf{I}_m \\ \mathbf{I}_{n_{\text{lat}}} & \mathbf{O} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{s}_{\text{lat}} \\ \mathbf{k} \end{bmatrix} + \begin{bmatrix} \mathbf{e} + \mathbf{A}_{\text{enu}}\mathbf{s}_{\text{enu}} \\ -\mathbf{s}_{\text{lat}} \end{bmatrix} \quad (3)$$

for some $\mathbf{k} \in \mathbb{Z}^m$. Denote

$$\mathbf{B} = \begin{bmatrix} \mathbf{A}_{\text{lat}} & q\mathbf{I}_m \\ \mathbf{I}_{n_{\text{lat}}} & \mathbf{O} \end{bmatrix}. \quad (4)$$

Let $L(\mathbf{B})$ be the lattice with basis $\mathbf{B}$. For each partition of I , we guess that all components indexed by I_{enu} are 0, i.e., $\mathbf{s}_{\text{enu}} = \mathbf{0}$. Then we use BKZ- β to obtain a reduced basis $\mathbf{c}_1, \dots, \mathbf{c}_n$ of L^* . Let $t < n$ be a positive integer, and let L_t be the sublattice with basis $\mathbf{c}_1, \dots, \mathbf{c}_t$. Using discrete Gaussian sampling, we obtain N vectors $\mathbf{u}_1, \dots, \mathbf{u}_N$ independently distributed according to $D_{L_t, \mathbf{s}}$. Then we use the FFT algorithm to compute the following score function for all $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$:

$$f(\tilde{\mathbf{s}}_{\text{fft}}) := \frac{1}{N} \sum_{j=1}^N \cos(2\pi \langle (\mathbf{b} - \mathbf{A}_{\text{fft}}\tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0}), \mathbf{u}_j \rangle), \quad (5)$$

where $s > 0$ is a real parameter to be determined, and its value will be analyzed in detail later. In the later analysis, we will give a threshold $T \in \mathbb{R}$. If $\max_{\tilde{s}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}} f(\tilde{s}_{\text{fft}}) \geq T$, then $\mathbf{s}_{\text{enu}} = \mathbf{0}$ with high probability; otherwise, $\mathbf{s}_{\text{enu}} \neq \mathbf{0}$ with high probability. Below we adopt the same strategy as in [6]: fix I_{lat} , uniformly and randomly choose n_{enu} elements from $I \setminus I_{\text{lat}}$ as I_{enu} , and the remaining ones as I_{fft} . We perform the above attack process, repeating until in some attack the maximum value of the score function is greater than or equal to the threshold T . Note that the advantage of fixing I_{lat} is that the BKZ algorithm and the N independent discrete Gaussian samplings only need to be performed once, and the resulting vectors are applicable to the subsequent loop [6]. The pseudocode of our attack algorithm is given below:

Algorithm 1 The dual attack to solve LWE using projection

Require: A sample $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ produced by an $\text{LWE}_{n,m,q,\chi_s,\chi_e}$ oracle.

Parameters: some positive integers $R, T, N, \beta, s, t, n_{\text{enu}}, n_{\text{fft}}, n_{\text{lat}}$.

Ensure: The secret vector $\mathbf{s}$.

- 1: Choose $I_{\text{lat}} \subseteq \{1, \dots, n\}$ such that $|I_{\text{lat}}| = n_{\text{lat}}$;
 - 2: $\mathbf{A}_{\text{lat}} \leftarrow$ select the columns of $\mathbf{A}$ indexed by I_{lat} ;
 - 3: Use BKZ with block size β on the lattice $L^*(\mathbf{B})$, where $\mathbf{B}$ is as in equation (4), to produce a reduced basis $\mathbf{c}_1, \dots, \mathbf{c}_n$ of $L^*(\mathbf{B})$;
 - 4: Use the discrete Gaussian sampling algorithm to obtain N vectors $\mathbf{u}_1, \dots, \mathbf{u}_N$ independently distributed according to $D_{L_t, s}$, where L_t is the sublattice of $L^*(\mathbf{B})$ with basis $\mathbf{c}_1, \dots, \mathbf{c}_t$.
 - 5: **for** R times **do**
 - 6: Choose a partition $I_{\text{enu}} \cup I_{\text{fft}}$ of $\{1, \dots, n\} \setminus I_{\text{lat}}$ with $|I_{\text{enu}}| = n_{\text{enu}}$ and $|I_{\text{fft}}| = n_{\text{fft}}$;
 - 7: $\mathbf{A}_{\text{enu}}, \mathbf{A}_{\text{fft}} \leftarrow$ select the columns of $\mathbf{A}$ indexed by I_{enu} and I_{fft} , respectively;
 - 8: Compute $f(\tilde{s}_{\text{fft}})$ for all $\tilde{s}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$, where f is defined in (5);
 - 9: $V \leftarrow \max_{\tilde{s}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}} f(\tilde{s}_{\text{fft}})$;
 - 10: **if** $V \geq T$ **then**
 - 11: $\mathbf{s}_{\text{enu}} \leftarrow \mathbf{0}$;
 - 12: $(\mathbf{s}_{\text{fft}}, \mathbf{s}_{\text{lat}}) \leftarrow \text{SUB\_LWE\_SOLVER}([\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{lat}}], \mathbf{b})$;
 - 13: **return** $(\mathbf{s}_{\text{enu}}, \mathbf{s}_{\text{fft}}, \mathbf{s}_{\text{lat}}, I_{\text{enu}}, I_{\text{fft}}, I_{\text{lat}})$;
 - 14: **end if**
 - 15: **end for**
-

We will prove that, with a suitable choice of parameters, with high probability, when $V \geq T$, we have $\mathbf{s}_{\text{enu}} = \mathbf{0}$. In this case, $\text{SUB\_LWE\_SOLVER}([\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{lat}}], \mathbf{b})$ solves an LWE problem of dimension $n_{\text{fft}} + n_{\text{lat}}$, i.e.,

$$\mathbf{b} = [\mathbf{A}_{\text{fft}}, \mathbf{A}_{\text{lat}}] \cdot \begin{bmatrix} \mathbf{s}_{\text{fft}} \\ \mathbf{s}_{\text{lat}} \end{bmatrix} + \mathbf{e} \bmod q,$$

and the running time of this subroutine is much smaller than the total time since $n_{\text{fft}} + n_{\text{lat}} < n$.

4.2 Correctness of the Algorithm

In this subsection, we prove under what conditions Algorithm 1 can recover the correct secret $\mathbf{s}$. Without loss of generality, assume $\mathbf{e} \neq \mathbf{0}$. We first prove that when $m \geq n$, the probability that $\mathbf{A}$ has full rank over $\mathbb{Z}_q$ is close to 1, and that when $n_{\text{enu}} + n_{\text{fft}} < m$, the probability that $\mathbf{e}$ can be expressed as a linear combination of the columns of $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ is close to 0.

Lemma 8. *Let $\mathbf{A}$ be chosen uniformly at random from $\mathbb{Z}_q^{m \times n}$, let each component of $\mathbf{e}$ be chosen independently according to the centered binomial distribution with parameter α , and assume $\mathbf{e} \neq \mathbf{0}$. If $m \geq n$, then*

$$\Pr[\mathbf{A} \text{ has full rank over } \mathbb{Z}_q] \geq 1 - \frac{q^n - 1}{q^m(q - 1)}.$$

Moreover, let E be the event that $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ has full column rank over $\mathbb{Z}_q$ and that $\mathbf{e}$ cannot be expressed as a linear combination of the columns of $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ over $\mathbb{Z}_q$. Then

$$\Pr[E] \geq 1 - \frac{q^{n'} - 1}{q^{m-1}(q - 1)},$$

where $n' = n_{\text{enu}} + n_{\text{fft}}$.

Proof. Since $m \geq n$, we analyze the columns of $\mathbf{A}$. Let the columns of $\mathbf{A}$ be $\mathbf{a}_1, \dots, \mathbf{a}_n$. Then $\mathbf{A}$ has full rank over $\mathbb{Z}_q$ if and only if $\mathbf{a}_1, \dots, \mathbf{a}_n$ are linearly independent over $\mathbb{Z}_q$. The columns are independent and identically distributed, so we analyze them column by column. First, $\Pr[\mathbf{a}_1 \neq \mathbf{0}] = 1 - \frac{1}{q^m}$. For any fixed $\mathbf{v}_1 \in \mathbb{Z}_q^m \setminus \{\mathbf{0}\}$, we have

$$\begin{aligned} & \Pr[\mathbf{a}_1, \mathbf{a}_2 \text{ are linearly independent over } \mathbb{Z}_q \mid \mathbf{a}_1 = \mathbf{v}_1] \\ &= \Pr[\mathbf{a}_2 \not\equiv k\mathbf{v}_1 \pmod{q} \text{ for all } k \in \mathbb{Z}] = 1 - \frac{1}{q^{m-1}}, \end{aligned}$$

since there are q vectors in $\mathbb{Z}_q^m$ that are linearly dependent with $\mathbf{v}_1$. Similarly, once $\mathbf{a}_1 = \mathbf{v}_1 \neq \mathbf{0}$ and $\mathbf{a}_2 = \mathbf{v}_2 \neq \mathbf{0}$ are fixed, the probability that $\mathbf{a}_3$ is linearly independent of $\mathbf{a}_1$ and $\mathbf{a}_2$ is $1 - \frac{1}{q^{m-2}}$. By induction, we obtain

$$\begin{aligned} & \Pr[\mathbf{a}_1, \dots, \mathbf{a}_n \text{ are linearly independent over } \mathbb{Z}_q] \\ &= \prod_{i=0}^{n-1} \left(1 - \frac{1}{q^{m-i}}\right) \geq 1 - \sum_{i=0}^{n-1} \frac{1}{q^{m-i}} = 1 - \frac{q^n - 1}{q^m(q - 1)}. \end{aligned}$$

Fix any $\mathbf{e} \in \mathbb{Z}_q^m \setminus \{\mathbf{0}\}$. Let the columns of $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ be $\mathbf{t}_1, \dots, \mathbf{t}_{n'}$, where $n' = n_{\text{enu}} + n_{\text{fft}} < m$. Then $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ has full column rank over $\mathbb{Z}_q$ and $\mathbf{e}$ cannot be expressed as a linear combination of $\mathbf{t}_1, \dots, \mathbf{t}_{n'}$ if and only if $\mathbf{e}, \mathbf{t}_1, \dots, \mathbf{t}_{n'}$ are linearly independent over $\mathbb{Z}_q$. Using the above analysis, the probability of this event is

$$\prod_{i=1}^{n'} \left(1 - \frac{1}{q^{m-i}}\right) \geq 1 - \sum_{i=1}^{n'} \frac{1}{q^{m-i}} = 1 - \frac{q^{n'} - 1}{q^{m-1}(q - 1)}.$$

Since $\mathbf{e}$ and $\mathbf{A}$ are independently distributed, by the law of total probability, when $\mathbf{e} \leftarrow \chi_{\mathbf{e}}$ with $\mathbf{e} \neq \mathbf{0}$ and $\mathbf{A} \leftarrow \mathcal{U}(\mathbb{Z}_q^{m \times n})$, the probability that $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ has full column rank over $\mathbb{Z}_q$ and $\mathbf{e}$ cannot be expressed as a linear combination over $\mathbb{Z}_q$ of the columns of $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ is at least $1 - \frac{q^{n'} - 1}{q^{m-1}(q-1)}$. $\square$

In practical attacks, one usually takes $m > n$ and $n_{\text{enu}} + n_{\text{fft}}$ much smaller than m , so both probabilities in Lemma 8 are exponentially close to 1. Hence, without loss of generality, in the following analysis we assume that $\mathbf{A}$ has full rank over $\mathbb{Z}_q$, that $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ has full column rank over $\mathbb{Z}_q$, and that $\mathbf{e}$ cannot be expressed as a linear combination of the columns of $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ over $\mathbb{Z}_q$.

Recall that L_t is the lattice in Step 4 of Algorithm 1. Let $V = \text{span}(L_t)$ be the subspace spanned by $\mathbf{c}_1, \dots, \mathbf{c}_t$, and let $\pi : \mathbb{R}^{m+n_{\text{lat}}} \rightarrow V$ denote the orthogonal projection of $\mathbb{R}^{m+n_{\text{lat}}}$ onto V . We now analyze the cases where the guess is correct and incorrect.

When the guess is correct, i.e., $\mathbf{s}_{\text{enu}} = \mathbf{0}$, it follows from (3) that

$$\begin{bmatrix} \mathbf{b} - \mathbf{A}_{\text{fft}} \mathbf{s}_{\text{fft}} \\ \mathbf{0} \end{bmatrix} \in L(\mathbf{B}) + \begin{bmatrix} \mathbf{e} \\ -\mathbf{s}_{\text{lat}} \end{bmatrix}.$$

Projecting both sides orthogonally onto V yields

$$\pi \left(\begin{bmatrix} \mathbf{b} - \mathbf{A}_{\text{fft}} \mathbf{s}_{\text{fft}} \\ \mathbf{0} \end{bmatrix} \right) \in \pi(L(\mathbf{B})) + \pi \left(\begin{bmatrix} \mathbf{e} \\ -\mathbf{s}_{\text{lat}} \end{bmatrix} \right).$$

Therefore,

$$\text{dist}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \mathbf{s}_{\text{fft}}, \mathbf{0})), \pi(L(\mathbf{B}))) \leq \|\pi((\mathbf{e}, -\mathbf{s}_{\text{lat}}))\|. \quad (6)$$

When the guess is incorrect, i.e., $\mathbf{s}_{\text{enu}} \neq \mathbf{0}$, let L' be the lattice generated by the columns of

$$\mathbf{B}' = \begin{bmatrix} \mathbf{A}_{\text{enu}} & \mathbf{A}_{\text{fft}} & \mathbf{A}_{\text{lat}} & q\mathbf{I}_m \\ \mathbf{O} & \mathbf{O} & \mathbf{I}_{\text{lat}} & \mathbf{O} \end{bmatrix}.$$

Write $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} + q\mathbf{k}$ for some $\mathbf{k} \in \mathbb{Z}^m$. Then for any $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ and any $(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^{n_{\text{lat}}+m}$, we have

$$\begin{aligned} & \begin{bmatrix} \mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{A}_{\text{lat}} & q\mathbf{I}_m \\ \mathbf{I}_{\text{lat}} & \mathbf{O} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{A}_{\text{enu}} \mathbf{s}_{\text{enu}} + \mathbf{A}_{\text{fft}} (\mathbf{s}_{\text{fft}} - \tilde{\mathbf{s}}_{\text{fft}}) + \mathbf{e} + q\mathbf{k} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{A}_{\text{lat}} & q\mathbf{I}_m \\ \mathbf{I}_{\text{lat}} & \mathbf{O} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{e} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{A}_{\text{enu}} & \mathbf{A}_{\text{fft}} & \mathbf{A}_{\text{lat}} & q\mathbf{I}_m \\ \mathbf{O} & \mathbf{O} & \mathbf{I}_{\text{lat}} & \mathbf{O} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{s}_{\text{enu}} \\ \mathbf{s}_{\text{fft}} - \tilde{\mathbf{s}}_{\text{fft}} \\ \mathbf{x} \\ \mathbf{y} + \mathbf{k} \end{bmatrix}. \end{aligned}$$

Denote the second term on the right-hand side of the last equation by $\mathbf{v}(\mathbf{x}, \mathbf{y})$. Then it is easy to see that $\mathbf{v}(\mathbf{x}, \mathbf{y}) = \mathbf{0}$ if and only if $\mathbf{x} = \mathbf{0}$ and $\mathbf{A}_{\text{enu}} \mathbf{s}_{\text{enu}} +$

$\mathbf{A}_{\text{fft}}(\mathbf{s}_{\text{fft}} - \tilde{\mathbf{s}}_{\text{fft}}) \equiv \mathbf{0} \pmod{q}$, which is impossible since $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ has full column rank over $\mathbb{Z}_q$ and $\mathbf{s}_{\text{enu}} \neq \mathbf{0}$. Therefore, $\mathbf{v}(\mathbf{x}, \mathbf{y}) \in L' \setminus \{\mathbf{0}\}$. Moreover, since $\mathbf{e}$ cannot be expressed as a linear combination of the columns of $[\mathbf{A}_{\text{enu}} \mathbf{A}_{\text{fft}}]$ over $\mathbb{Z}_q$, we have $(\mathbf{e}, \mathbf{0}) + \mathbf{v}(\mathbf{x}, \mathbf{y}) \neq \mathbf{0}$.

By the arbitrariness of $(\mathbf{x}, \mathbf{y})$, it follows that $(\mathbf{b} - \mathbf{A}_{\text{fft}}\tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0}) \notin L(\mathbf{B})$. Projecting both sides of the above equation onto V yields

$$\pi\left(\begin{bmatrix} \mathbf{b} - \mathbf{A}_{\text{fft}}\tilde{\mathbf{s}}_{\text{fft}} \\ \mathbf{0} \end{bmatrix}\right) + \pi\left(\begin{bmatrix} \mathbf{A}_{\text{lat}} & q\mathbf{I}_m \\ \mathbf{I}_{\text{lat}} & \mathbf{O} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}\right) = \pi\left(\begin{bmatrix} \mathbf{e} \\ \mathbf{0} \end{bmatrix}\right) + \pi(\mathbf{v}(\mathbf{x}, \mathbf{y})).$$

Without loss of generality, we assume that $\pi(\mathbf{v}(\mathbf{x}, \mathbf{y})) \neq \mathbf{0}$ for all $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ and $(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^{n_{\text{lat}}+m}$. Under this condition, we have

$$\begin{aligned} \text{dist}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}}\tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})), \pi(L(\mathbf{B}))) &= \min_{(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^{n_{\text{lat}}+m}} \|\pi((\mathbf{e}, \mathbf{0})) + \pi(\mathbf{v}(\mathbf{x}, \mathbf{y}))\| \\ &\geq \min_{(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^{n_{\text{lat}}+m}} \|\pi(\mathbf{v}(\mathbf{x}, \mathbf{y}))\| - \|\pi((\mathbf{e}, \mathbf{0}))\| \\ &\geq \lambda_1(\pi(L')) - \|\pi((\mathbf{e}, \mathbf{0}))\|. \end{aligned} \quad (7)$$

We note that if there exist $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ and $(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^{n_{\text{lat}}+m}$ such that $\pi(\mathbf{v}(\mathbf{x}, \mathbf{y})) = \mathbf{0}$, then the attacker can instead solve a uSVP or BDD instance on a lattice $\pi(L')$ of rank only t . Specifically, this degenerate condition yields

$$\pi\left(\begin{bmatrix} \mathbf{b} - \mathbf{A}_{\text{fft}}\tilde{\mathbf{s}}_{\text{fft}} \\ \mathbf{0} \end{bmatrix}\right) + \pi\left(\begin{bmatrix} \mathbf{A}_{\text{lat}} & q\mathbf{I}_m \\ \mathbf{I}_{\text{lat}} & \mathbf{O} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}\right) = \pi\left(\begin{bmatrix} \mathbf{e} \\ \mathbf{0} \end{bmatrix}\right),$$

or equivalently,

$$\pi\left(\begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix}\right) + \pi\left(\begin{bmatrix} \mathbf{A}_{\text{fft}} & \mathbf{A}_{\text{lat}} & q\mathbf{I}_m \\ \mathbf{O} & \mathbf{I}_{\text{lat}} & \mathbf{O} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{s}}_{\text{fft}} \\ \mathbf{x} \\ \mathbf{y} \end{bmatrix}\right) = \pi\left(\begin{bmatrix} \mathbf{e} \\ \mathbf{0} \end{bmatrix}\right).$$

Under the GSA assumption and the central limit theorem, $\pi((\mathbf{e}, \mathbf{0}))$ is nonzero with overwhelming probability. If $\pi((\mathbf{b}, \mathbf{0})) \in \pi(L')$ (equivalently, $\pi((\mathbf{e}, \mathbf{0})) \in \pi(L')$), then $\pi(L')$ contains a nonzero vector $\pi((\mathbf{e}, \mathbf{0}))$ of length at most $\|\mathbf{e}\|$, which is an unusually short vector in the lattice for typical parameters; solving uSVP on $\pi(L')$ recovers this vector. If $\pi((\mathbf{b}, \mathbf{0})) \notin \pi(L')$, then the attacker faces a BDD instance $(\pi((\mathbf{b}, \mathbf{0})), \pi(L'))$ on $\pi(L')$, whose target vector is at distance $\|\pi((\mathbf{e}, \mathbf{0}))\|$ from the lattice and can be handled by a BDD solver. In either case, the attacker only needs to handle a lattice $\pi(L')$ of rank t , which is typically much smaller than the original lattice dimension. Clearly, the computational cost of solving a rank- t lattice problem does not exceed that of Algorithm 1 on the original high-dimensional lattice. Therefore, such degenerate instances do not increase the worst-case running time of the attack algorithm, and we can safely assume in the subsequent analysis that this degenerate case does not occur.

Since the function $\frac{\rho_{1/s}(\mathbf{t}+L)}{\rho_{1/s}(L)}$ reflects the distance between $\mathbf{t}$ and L , using Corollary 3 we obtain the following result:

Lemma 9. Assume that $\lambda_1(\pi(L')) - \|\pi((e, \mathbf{0}))\| \geq \frac{c}{s} \sqrt{\frac{t}{2\pi}}$ and $\lambda_1(L(\pi(\mathbf{B}))) \geq \frac{kc}{s} \sqrt{\frac{t}{2\pi}}$, where $k \geq 1.04$, $c \geq 1$, and $t \geq 500$. Then, when $\mathbf{s}_{\text{enu}} = \mathbf{0}$, we have

$$\frac{\rho_{1/s}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \mathbf{s}_{\text{fft}}, \mathbf{0})) + \pi(L(\mathbf{B})))}{\rho_{1/s}(\pi(L(\mathbf{B})))} \geq \rho_{1/s}(\pi((e, -\mathbf{s}_{\text{lat}}))). \quad (8)$$

When $\mathbf{s}_{\text{enu}} \neq \mathbf{0}$, for any $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ we have

$$\frac{\rho_{1/s}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})) + \pi(L(\mathbf{B})))}{\rho_{1/s}(\pi(L(\mathbf{B})))} \leq 2 \left(e^{1-c^2} \right)^{\frac{t}{2}}. \quad (9)$$

Proof. This follows directly from (6), (7), Lemma 7, and Corollary 3. $\square$

Using the Poisson summation formula and Hoeffding's inequality, we can estimate $\frac{\rho_{1/s}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})) + \pi(L(\mathbf{B})))}{\rho_{1/s}(\pi(L(\mathbf{B})))}$. Specifically, we have the following result.

Lemma 10. Let $s, \delta > 0$ and $N \in \mathbb{N}$. Then with probability at least $1 - 2Rq^{n_{\text{fft}}} e^{-\frac{N\delta^2}{2}}$ over the choice of $U = \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$ drawn independently from $D_{L_t, s}$, we have

$$\left| f(\tilde{\mathbf{s}}_{\text{fft}}) - \frac{\rho_{1/s}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})) + \pi(L(\mathbf{B})))}{\rho_{1/s}(\pi(L(\mathbf{B})))} \right| \leq \delta$$

for all choices of I_{enu} and I_{fft} and all $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ in the R iterations of Algorithm 1. Here $f(\tilde{\mathbf{s}}_{\text{fft}})$ is defined in (5).

Proof. By Lemma 1, we have $L_t^* = \pi(L(\mathbf{B}))$. Then, by the Poisson summation formula,

$$\begin{aligned} \frac{\rho_{1/s}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})) + \pi(L(\mathbf{B})))}{\rho_{1/s}(\pi(L(\mathbf{B})))} &= \frac{\sum_{\mathbf{v} \in L_t} e^{2\pi i \langle \pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})), \mathbf{v} \rangle} \rho_s(\mathbf{v})}{\rho_s(L_t)} \\ &= \frac{\sum_{\mathbf{v} \in L_t} \cos(2\pi \langle \pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})), \mathbf{v} \rangle) \rho_s(\mathbf{v})}{\rho_s(L_t)} \\ &= \frac{\sum_{\mathbf{v} \in L_t} \cos(2\pi \langle (\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0}), \mathbf{v} \rangle) \rho_s(\mathbf{v})}{\rho_s(L_t)}. \end{aligned}$$

The last equality holds because $(\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0}) - \pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0}))$ is orthogonal to the subspace V spanned by L_t . Note that

$$\frac{\sum_{\mathbf{v} \in L_t} \cos(2\pi \langle (\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0}), \mathbf{v} \rangle) \rho_s(\mathbf{v})}{\rho_s(L_t)} = \mathbb{E}_{\mathbf{v} \leftarrow D_{L_t, s}} [\cos(2\pi \langle (\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0}), \mathbf{v} \rangle)].$$

For any $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ in any of the R iterations of Algorithm 1, using Hoeffding's inequality (Theorem 1) we obtain

$$\Pr \left[\left| f(\tilde{\mathbf{s}}_{\text{fft}}) - \frac{\rho_{1/s}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})) + \pi(L(\mathbf{B})))}{\rho_{1/s}(\pi(L(\mathbf{B})))} \right| \geq \delta \right] \leq 2e^{-\frac{N\delta^2}{2}},$$

since $\cos(2\pi \langle (\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0}), \mathbf{v} \rangle) \in [-1, 1]$. Finally, applying the union bound proves the claim. $\square$

The following result gives a lower bound on the probability that, in the R iterations, at least one iteration satisfies $\mathbf{s}_{\text{enu}} = \mathbf{0}$.

Lemma 11 ([6, Lemma 3.2]). *Let each component of $\mathbf{s}$ be independently distributed according to the centered binomial distribution with parameter α , where α is a nonzero positive integer. The probability that Algorithm 1 has at least one iteration among the R iterations with $\mathbf{s}_{\text{enu}} = \mathbf{0}$ is*

$$\eta := 1 - \sum_{k=0}^{n_{\text{enu}}+n_{\text{fft}}} \left(1 - \frac{\binom{k}{n_{\text{enu}}}}{\binom{n_{\text{enu}}+n_{\text{fft}}}{n_{\text{enu}}}}\right)^R \binom{n_{\text{enu}}+n_{\text{fft}}}{k} p_0^k (1-p_0)^{n_{\text{enu}}+n_{\text{fft}}-k}, \quad (10)$$

where $p_0 := 2^{-2\alpha} \binom{2\alpha}{\alpha}$.

With the above preparations, we are now ready to prove the correctness of Algorithm 1.

Theorem 3 (Correctness of Algorithm 1). *Let $s, \delta, T > 0$ and $N \in \mathbb{N}$. Assume that $\lambda_1(\pi(L')) - \|\pi((\mathbf{e}, \mathbf{0}))\| \geq \frac{c}{s} \sqrt{\frac{t}{2\pi}}$ and $\lambda_1(\pi(L(\mathbf{B}))) \geq \frac{kc}{s} \sqrt{\frac{t}{2\pi}}$, where $k \geq 1.04$, $c \geq 1$, and $t \geq 500$. Suppose that SUB_LWE_SOLVER outputs the correct $\mathbf{s}_{\text{fft}}$ and $\mathbf{s}_{\text{lat}}$ under the condition $\mathbf{s}_{\text{enu}} = \mathbf{0}$, and suppose that*

$$\rho_{1/s}(\pi((\mathbf{e}, -\mathbf{s}_{\text{lat}}))) - \delta > T > 2 \left(e^{1-c^2} \right)^{\frac{t}{2}} + \delta. \quad (11)$$

Then, with probability at least $\eta(1 - 2Rq^{n_{\text{fft}}} e^{-\frac{N\delta^2}{2}})$, Algorithm 1 eventually outputs the correct secret $\mathbf{s}$.

Proof. By Lemma 9, in the R iterations of Algorithm 1, when $\mathbf{s}_{\text{enu}} = \mathbf{0}$, there exists $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ such that

$$\frac{\rho_{1/s}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})) + \pi(L(\mathbf{B})))}{\rho_{1/s}(\pi(L(\mathbf{B})))} \geq \rho_{1/s}(\pi((\mathbf{e}, -\mathbf{s}_{\text{lat}}))).$$

When $\mathbf{s}_{\text{enu}} \neq \mathbf{0}$, for any $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ we have

$$\frac{\rho_{1/s}(\pi((\mathbf{b} - \mathbf{A}_{\text{fft}} \tilde{\mathbf{s}}_{\text{fft}}, \mathbf{0})) + \pi(L(\mathbf{B})))}{\rho_{1/s}(\pi(L(\mathbf{B})))} \leq 2 \left(e^{1-c^2} \right)^{\frac{t}{2}}.$$

Then, by Lemma 10 and the triangle inequality, with probability at least $1 - 2Rq^{n_{\text{fft}}} e^{-\frac{N\delta^2}{2}}$, in the R iterations of Algorithm 1, when $\mathbf{s}_{\text{enu}} = \mathbf{0}$ we have

$$f(\mathbf{s}_{\text{fft}}) \geq \rho_{1/s}(\pi((\mathbf{e}, -\mathbf{s}_{\text{lat}}))) - \delta > T.$$

When $\mathbf{s}_{\text{enu}} \neq \mathbf{0}$, for any $\tilde{\mathbf{s}}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ we have

$$f(\tilde{\mathbf{s}}_{\text{fft}}) \leq 2 \left(e^{1-c^2} \right)^{\frac{t}{2}} + \delta < T.$$

Therefore, as long as in the R iterations of Algorithm 1 there is at least one iteration with $\mathbf{s}_{\text{enu}} = \mathbf{0}$, Steps 10 to 14 of Algorithm 1 guarantee that the correct $\mathbf{s}$ is eventually output. By Lemma 11, the probability that in the R iterations there is at least one iteration with $\mathbf{s}_{\text{enu}} = \mathbf{0}$ is at least η . Since the choices of $\mathbf{s}$ and I_{enu} are independent of the choice of U , the probability that Algorithm 1 outputs the correct $\mathbf{s}$ is at least $\eta(1 - 2Rq^{n_{\text{fft}}}e^{-\frac{N\delta^2}{2}})$. $\square$

4.3 Parameter Selection for Kyber

In this subsection, we analyze under which concrete parameter choices Algorithm 1 outputs the correct $\mathbf{s}$ with constant probability. First, we choose a suitable R such that $\eta \geq 0.5$. Later, we will choose suitable N and δ such that $1 - 2Rq^{n_{\text{fft}}}e^{-\frac{N\delta^2}{2}}$ tends to 1, so that the success probability of Algorithm 1 is approximately 0.5. Note that in the Kyber scheme, both $\mathbf{s}$ and $\mathbf{e}$ follow the centered binomial distribution; denote the standard deviation of this distribution by σ .

By Hoeffding's inequality, for any $\tau > 0$,

$$\Pr [\|\mathbf{e}\| > \sqrt{m}\sigma + \tau] \leq 2^{-\Omega(\tau^2)},$$

since $\mathbb{E}[\|\mathbf{e}\|] \leq \sqrt{\mathbb{E}[\|\mathbf{e}\|^2]} = \sqrt{m} \cdot \sigma$. A similar conclusion holds for $(\mathbf{e}, \mathbf{s}_{\text{lat}})$. Therefore, with high probability,

$$\|\mathbf{e}\| = \sqrt{m} \cdot \sigma + O(1), \quad \|(\mathbf{e}, \mathbf{s}_{\text{lat}})\| = \sqrt{m + n_{\text{lat}}} \cdot \sigma + O(1).$$

To simplify the analysis, we ignore the lower-order terms and directly use the above typical lengths.

Note that V is the subspace spanned by the first t vectors of the reduced basis obtained by applying BKZ to a basis of $L^*(\mathbf{B})$. We adopt the standard random direction heuristic in lattice cryptanalysis: the reduced basis vectors behave as random directions relative to the secret/noise vectors. Under this heuristic,

$$\mathbb{E} [\|\pi((\mathbf{e}, \mathbf{s}_{\text{lat}}))\|^2] = t\sigma^2, \quad \mathbb{E} [\|\pi((\mathbf{e}, \mathbf{0}))\|^2] = \frac{tm}{m + n_{\text{lat}}}\sigma^2.$$

Hence we use the typical lengths

$$\|\pi((\mathbf{e}, \mathbf{s}_{\text{lat}}))\| \approx \sqrt{t} \cdot \sigma, \quad \|\pi((\mathbf{e}, \mathbf{0}))\| \approx \sqrt{\frac{tm}{m + n_{\text{lat}}}} \cdot \sigma$$

to estimate the attack complexity, ignoring lower-order statistical fluctuations. This treatment is consistent with common heuristics in existing lattice cryptanalysis.

We choose δ sufficiently smaller than the lower bound $\rho_{1/s}(\pi((\mathbf{e}, \mathbf{s}_{\text{lat}})))$ of the score function when the guess is correct. Specifically, we take

$$\delta = \frac{1}{10}\rho_{1/s}(\pi((\mathbf{e}, \mathbf{s}_{\text{lat}}))) = \frac{1}{10}e^{-\pi s^2 t \sigma^2}. \quad (12)$$

We choose T suitably smaller than $\rho_{1/s}(\pi((\mathbf{e}, -\mathbf{s}_{\text{lat}}))) - \delta$. For instance, we take

$$T = \frac{4}{5} \rho_{1/s}(\pi((\mathbf{e}, -\mathbf{s}_{\text{lat}}))) = \frac{4}{5} e^{-\pi s^2 t \sigma^2}. \quad (13)$$

We choose N sufficiently large so that $1 - 2Rq^{n_{\text{fft}}} e^{-\frac{N\delta^2}{2}}$ tends to 1. Therefore, let

$$N = \frac{2(\text{poly}(n) + n_{\text{fft}} \log q + \log R)}{\delta^2}. \quad (14)$$

Then

$$1 - 2Rq^{n_{\text{fft}}} e^{-\frac{N\delta^2}{2}} = 1 - 2e^{-\text{poly}(n)},$$

which is exponentially close to 1.

Finally, by Theorem 3, we additionally require that

$$\lambda_1(\pi(L')) - \|\pi((\mathbf{e}, \mathbf{0}))\| \geq \frac{c}{s} \sqrt{\frac{t}{2\pi}}$$

and

$$\lambda_1(\pi(L(\mathbf{B}))) \geq \frac{kc}{s} \sqrt{\frac{t}{2\pi}},$$

where $k \geq 1.04$, $c \geq 1$, and $t \geq 500$, and that

$$T > 2 \left(e^{1-c^2} \right)^{\frac{t}{2}} + \delta.$$

To verify these conditions, we first compute the determinants of $L(\mathbf{B})$ and L' .

Lemma 12.

$$\det(L(\mathbf{B})) = q^m, \quad \det(L') = q^{m-n_{\text{enu}}-n_{\text{fft}}}.$$

Proof. Since $\mathbf{B}$ itself is a lattice basis of $L(\mathbf{B})$, we have $\det(L(\mathbf{B})) = \det(\mathbf{B}) = q^m$. By the definition of L' , the lattice $L(\mathbf{B})$ is a sublattice of L' . Take any $\mathbf{v} \in L'$. Then by definition there exist $\mathbf{x}_1 \in \mathbb{Z}^{n_{\text{enu}}}$, $\mathbf{x}_2 \in \mathbb{Z}^{n_{\text{fft}}}$, $\mathbf{x}_3 \in \mathbb{Z}^{n_{\text{lat}}}$, and $\mathbf{x}_4 \in \mathbb{Z}^m$ such that

$$\mathbf{v} = \begin{bmatrix} \mathbf{A}_{\text{enu}}\mathbf{x}_1 + \mathbf{A}_{\text{fft}}\mathbf{x}_2 \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{A}_{\text{lat}}\mathbf{x}_3 + q\mathbf{x}_4 \\ \mathbf{x}_3 \end{bmatrix},$$

where $(\mathbf{A}_{\text{lat}}\mathbf{x}_3 + q\mathbf{x}_4, \mathbf{x}_3) \in L(\mathbf{B})$. Hence $\mathbf{v} \notin L(\mathbf{B})$ if and only if $(\mathbf{A}_{\text{enu}}\mathbf{x}_1 + \mathbf{A}_{\text{fft}}\mathbf{x}_2, \mathbf{0}) \notin L(\mathbf{B})$. This is equivalent to $\mathbf{A}_{\text{enu}}\mathbf{x}_1 + \mathbf{A}_{\text{fft}}\mathbf{x}_2 \not\equiv \mathbf{0} \pmod{q}$. Since $[\mathbf{A}_{\text{enu}} \ \mathbf{A}_{\text{fft}}]$ has full column rank over $\mathbb{Z}_q$, the map $(\mathbf{x}_1, \mathbf{x}_2) \mapsto \mathbf{A}_{\text{enu}}\mathbf{x}_1 + \mathbf{A}_{\text{fft}}\mathbf{x}_2 \pmod{q}$ is injective, and therefore $[L' : L(\mathbf{B})] = q^{n_{\text{enu}}+n_{\text{fft}}}$. Consequently,

$$\det(L') = \frac{\det(L(\mathbf{B}))}{[L' : L(\mathbf{B})]} = q^{m-n_{\text{enu}}-n_{\text{fft}}}.$$

□

By Lemma 1, we have $\pi(L(\mathbf{B})) = L_t^*$. Then, by Lemma 12, under the GH and GSA assumptions,

$$\det(L_t) = \delta(\beta)^{t(n_{\text{lat}}+m-t)} \det(L^*(\mathbf{B}))^{\frac{t}{n_{\text{lat}}+m}} = \delta(\beta)^{t(n_{\text{lat}}+m-t)} q^{-\frac{mt}{n_{\text{lat}}+m}},$$

where $\delta(\beta)$ is the root Hermite factor. Hence

$$\det(\pi(L(\mathbf{B}))) = \frac{1}{\det(L_t)} = \delta(\beta)^{-t(n_{\text{lat}}+m-t)} q^{\frac{mt}{n_{\text{lat}}+m}}.$$

We wish to estimate $\det(\pi(L'))$ under the same heuristic. Since $L(\mathbf{B})$ is a sublattice of L' of index $q^{n_{\text{enu}}+n_{\text{fft}}}$, under the GSA and the random direction heuristic, the index $q^{n_{\text{enu}}+n_{\text{fft}}}$ shrinks by a factor of $q^{-t(n_{\text{enu}}+n_{\text{fft}})/(n_{\text{lat}}+m)}$ in the direction of V . Therefore, $\det(\pi(L'))$ can be estimated as

$$\det(\pi(L')) \approx q^{-t(n_{\text{enu}}+n_{\text{fft}})/(n_{\text{lat}}+m)} \det(\pi(L(\mathbf{B}))) = \delta(\beta)^{-t(n_{\text{lat}}+m-t)} q^{\frac{t(m-n_{\text{enu}}-n_{\text{lat}})}{n_{\text{lat}}+m}}.$$

Using GH, we obtain

$$\lambda_1(\pi(L(\mathbf{B}))) \approx \sqrt{\frac{t}{2\pi e}} \det(L(\mathbf{B}))^{1/t} = \sqrt{\frac{t}{2\pi e}} \delta(\beta)^{-(n_{\text{lat}}+m-t)} q^{\frac{m}{n_{\text{lat}}+m}}$$

and

$$\lambda_1(\pi(L')) \approx \sqrt{\frac{t}{2\pi e}} \det(\pi(L'))^{1/t} \approx \sqrt{\frac{t}{2\pi e}} \delta(\beta)^{-(n_{\text{lat}}+m-t)} q^{\frac{m-n_{\text{enu}}-n_{\text{lat}}}{n_{\text{lat}}+m}}.$$

Under the above analysis, the conditions $\lambda_1(\pi(L')) - \|\pi((\mathbf{e}, \mathbf{0}))\| \geq \frac{c}{s} \sqrt{\frac{t}{2\pi}}$,

$\lambda_1(\pi(L(\mathbf{B}))) \geq \frac{kc}{s} \sqrt{\frac{t}{2\pi}}$, and $T > 2 \left(e^{1-c^2} \right)^{\frac{t}{2}} + \delta$ are equivalent to

$$\sqrt{\frac{t}{2\pi e}} \delta(\beta)^{-(n_{\text{lat}}+m-t)} q^{\frac{m-n_{\text{enu}}-n_{\text{lat}}}{n_{\text{lat}}+m}} - \sqrt{\frac{tm}{m+n_{\text{lat}}}} \cdot \sigma \geq \frac{c}{s} \sqrt{\frac{t}{2\pi}}, \quad (15)$$

$$\sqrt{\frac{t}{2\pi e}} \delta(\beta)^{-(n_{\text{lat}}+m-t)} q^{\frac{m}{n_{\text{lat}}+m}} \geq \frac{kc}{s} \sqrt{\frac{t}{2\pi}}, \quad (16)$$

and

$$\frac{7}{10} e^{-\pi s^2 t \sigma^2} > 2 \left(e^{1-c^2} \right)^{\frac{t}{2}}, \quad (17)$$

respectively. The following result states the conditions on c and s under which (15), (16), and (17) hold.

Lemma 13. *When*

$$s \geq \max \left\{ \left(\frac{1.01}{\left(\sqrt{\frac{1}{e}} \delta(\beta)^{-(n_{\text{lat}}+m-t)} q^{\frac{m-n_{\text{enu}}-n_{\text{lat}}}{n_{\text{lat}}+m}} - \sqrt{\frac{2\pi m}{m+n_{\text{lat}}}} \cdot \sigma \right)^2 - 2\pi \sigma^2} \right)^{\frac{1}{2}}, \left(\frac{1.01}{\frac{1}{k^2 e} \delta(\beta)^{-2(n_{\text{lat}}+m-t)} q^{\frac{2m}{n_{\text{lat}}+m}} - 2\pi \sigma^2} \right)^{\frac{1}{2}} \right\} \quad (18)$$

the inequalities (15), (16), and (17) hold.

Proof. We first prove that when $c \geq \sqrt{2\pi s^2 \sigma^2 + 1.01}$, (17) holds. In this case,

$$2 \left(e^{1-c^2} \right)^{\frac{t}{2}} \leq 2 \left(e^{1-(2\pi s^2 \sigma^2 + 1.01)} \right)^{\frac{t}{2}} = e^{-\pi s^2 t \sigma^2} (2e^{-0.005t}) < \frac{7}{10} e^{-\pi s^2 t \sigma^2}$$

since $t \geq 500$. It remains to prove that when $c = \sqrt{2\pi s^2 \sigma^2 + 1.01}$ and s satisfies the lower bound in (18), (15) and (16) hold. It suffices to substitute the lower bound on s into them respectively. $\square$

4.4 Complexity Analysis

In this subsection, we analyze the theoretical complexity of Algorithm 1, as well as the concrete complexity estimate of the attack under the specific parameters of Kyber. The theoretical complexity of Algorithm 1 is

$$T_{\text{BKZ}}(n_{\text{lat}} + m, \beta) + NT_{\text{sampling}}(D_{L_t, s}) + R(N \text{poly}(m, n) + T_{\text{FFT}}(\mathbb{Z}_q^{n_{\text{fft}}}).$$

Here T_{BKZ} is the complexity given in Lemma 5, and $T_{\text{sampling}}(D_{L_t, s})$ is the complexity of sampling one vector according to the distribution $D_{L_t, s}$, as given in Theorem 2. The term $R(N \text{poly}(m, n) + T_{\text{FFT}}(\mathbb{Z}_q^{n_{\text{fft}}}))$ is the complexity of computing $f(\tilde{s}_{\text{fft}})$ for all R iterations and for all $\tilde{s}_{\text{fft}} \in \mathbb{Z}_q^{n_{\text{fft}}}$ in each iteration, where $T_{\text{FFT}}(\mathbb{Z}_q^{n_{\text{fft}}})$ is given in Lemma 2. We modified the script provided in [23] to compute the corresponding complexity for all possible choices of the attack parameters, and output the parameters corresponding to the lowest complexity. The results are shown in the following tables.

Table 1. Attack parameters and evaluation results

Scheme	Costs(bits)	$\log_2(N)$	R	n_{enu}	n_{fft}	n_{lat}	m	t	η	s	β
Kyber512	211.9	152.9	3	1	16	495	690	560	0.649	0.134	650
Kyber768	317.1	232.0	7	2	25	741	810	810	0.608	0.171	1020
Kyber1024	434.3	300.5	7	2	35	987	1180	1090	0.620	0.169	1440

Table 2. Comparison between our method and [23]

Scheme	Our method	[23]
Kyber512	211.9	238
Kyber768	317.1	347
Kyber1024	434.3	478

It can be seen that our method achieves an improvement of 26.1–43.7 bits over [23]. However, it is worth noting that, although the improvement is significant, there is still a considerable gap compared with the heuristic dual attack. This gap

may be attributed to the additional constraints imposed by the rigorous analysis and the conservative parameter choices required to guarantee correctness, which inevitably lead to a loss in efficiency. Narrowing this gap while maintaining rigorous guarantees remains an interesting direction for future work.

5 Conclusion and Future Directions

We observe that Banaszczyk’s inequality can be substantially improved in certain regimes. More precisely, by imposing the additional condition that $\lambda_1(L)$ is sufficiently large, we derive a refined inequality. When applied to the dual attack, this refinement yields sharper conditions for attack success.

In addition to this refinement, we propose a new provable dual attack framework. Unlike previous approaches that perform discrete Gaussian sampling over the entire dual lattice, our method performs BKZ and discrete Gaussian sampling on a sublattice of the dual lattice. By exploiting the projection relation between this dual sublattice and the primal lattice, we project the LWE problem onto the span of the dual sublattice and solve it therein. This framework improves the attack efficiency by 26.1–43.7 bits over the original provable dual attack framework.

Nevertheless, despite this improvement, a considerable gap remains relative to the heuristic dual attack. This gap can be attributed to the additional constraints introduced by the rigorous analysis and the conservative parameter choices needed to guarantee correctness, both of which incur an inevitable efficiency loss. A promising direction for future work is to close this gap while preserving rigorous guarantees. One concrete avenue is to integrate modulus switching into our framework; this technique is also aimed at improving attack efficiency but has not been incorporated here due to space limitations.

References

1. Aggarwal, D., Dadush, D., Regev, O., Stephens-Davidowitz, N.: Solving the Shortest Vector Problem in 2^n Time Using Discrete Gaussian Sampling: Extended Abstract. In: Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing. pp. 733–742. STOC ’15, Association for Computing Machinery, New York, NY, USA (2015), <https://doi.org/10.1145/2746539.2746606>
2. Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen* **296**, 625–635 (1993)
3. Becker, A., Ducas, L., Gama, N., Laarhoven, T.: New directions in nearest neighbor searching with applications to lattice sieving. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 10–24. SODA ’16, Society for Industrial and Applied Mathematics, USA (2016)
4. Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehle, D.: CRYSTALS - Kyber: A CCA-Secure Module-Lattice-Based KEM. In: 2018 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 353–367 (2018). <https://doi.org/10.1109/EuroSP.2018.00032>

5. Brakerski, Z.: Fully Homomorphic Encryption without Modulus Switching from Classical GapSVP. In: Safavi-Naini, R., Canetti, R. (eds.) *Advances in Cryptology – CRYPTO 2012*. pp. 868–886. Springer Berlin Heidelberg, Berlin, Heidelberg (2012), https://doi.org/10.1007/978-3-642-32009-5_50
6. Carrier, K., Meyer-Hilfiger, C., Shen, Y., Tillich, J.P.: Assessing the impact of a variant of matzov’s dual attack on kyber. In: Tauman Kalai, Y., Kamara, S.F. (eds.) *Advances in Cryptology – CRYPTO 2025*. pp. 444–476. Springer Nature Switzerland, Cham (2025)
7. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic Encryption for Arithmetic of Approximate Numbers. In: Takagi, T., Peyrin, T. (eds.) *Advances in Cryptology – ASIACRYPT 2017*. pp. 409–437. Springer International Publishing, Cham (2017), https://doi.org/10.1007/978-3-319-70694-8_15
8. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: Fast Fully Homomorphic Encryption Over the Torus. *J. Cryptol.* **33**(1), 34–91 (2020), <https://doi.org/10.1007/s00145-019-09319-x>
9. Ducas, L.: Shortest vector from lattice sieving: A few dimensions for free. In: Nielsen, J.B., Rijmen, V. (eds.) *Advances in Cryptology – EUROCRYPT 2018*. pp. 125–145. Springer International Publishing, Cham (2018)
10. Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme. *Transactions on Cryptographic Hardware and Embedded Systems* **2018**, Issue 1, 238–268 (2018). <https://doi.org/10.13154/tches.v2018.i1.238-268>
11. Ducas, L., Pulles, L.N.: Does the Dual-Sieve Attack on Learning with Errors Even Work? In: Handschuh, H., Lysyanskaya, A. (eds.) *Advances in Cryptology – CRYPTO 2023*. pp. 37–69. Springer Nature Switzerland, Cham (2023), https://doi.org/10.1007/978-3-031-38548-3_2
12. Espitau, T., Fouque, P.A., Gérard, F., Röss, M., Takahashi, A., Tibouchi, M., Wallet, A., Yu, Y.: Mitaka: A Simpler, Parallelizable, Maskable Variant of Falcon. In: Dunkelman, O., Dziembowski, S. (eds.) *Advances in Cryptology – EUROCRYPT 2022*. pp. 222–253. Springer International Publishing, Cham (2022), https://doi.org/10.1007/978-3-031-07082-2_9
13. Espitau, T., Joux, A., Kharchenko, N.: On a Dual/Hybrid Approach to Small Secret LWE. In: Bhargavan, K., Oswald, E., Prabhakaran, M. (eds.) *Progress in Cryptology – INDOCRYPT 2020*. pp. 440–462. Springer International Publishing, Cham (2020), https://doi.org/10.1007/978-3-030-65277-7_20
14. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*. pp. 197–206. STOC ’08, Association for Computing Machinery, New York, NY, USA (2008), <https://doi.org/10.1145/1374376.1374407>
15. Guo, Q., Johansson, T.: Faster Dual Lattice Attacks for Solving LWE with Applications to CRYSTALS. In: Tibouchi, M., Wang, H. (eds.) *Advances in Cryptology – ASIACRYPT 2021*. pp. 33–62. Springer International Publishing, Cham (2021), https://doi.org/10.1007/978-3-030-92068-5_2
16. Klein, P.: Finding the closest lattice vector when it’s unusually close. In: *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms*. pp. 937–941. SODA ’00, Society for Industrial and Applied Mathematics, USA (2000)
17. Lyubashevsky, V., Nguyen, N.K., Plançon, M.: Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General. In: Dodis, Y., Shrimpton, T. (eds.) *Advances in Cryptology – CRYPTO 2022*. pp. 71–101. Springer Nature Switzerland, Cham (2022), https://doi.org/10.1007/978-3-031-15979-4_3

18. MATZOV: Report on the Security of LWE: Improved Dual Lattice Attack (2020), <https://doi.org/10.5281/zenodo.6493704>
19. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. In: 45th Annual IEEE Symposium on Foundations of Computer Science. pp. 372–381 (2004). <https://doi.org/10.1109/FOCS.2004.72>
20. Nguyen, P.Q., Vidick, T.: Sieve algorithms for the shortest vector problem are practical. *Journal of Mathematical Cryptology* **2**(2), 181–207 (2008)
21. Peikert, C.: An Efficient and Parallel Gaussian Sampler for Lattices. In: Rabin, T. (ed.) *Advances in Cryptology – CRYPTO 2010*. pp. 80–97. Springer Berlin Heidelberg, Berlin, Heidelberg (2010), https://doi.org/10.1007/978-3-642-14623-7_5
22. Peikert, C., Shiehian, S.: Noninteractive Zero Knowledge for NP from (Plain) Learning with Errors. In: Boldyreva, A., Micciancio, D. (eds.) *Advances in Cryptology – CRYPTO 2019*. pp. 89–114. Springer International Publishing, Cham (2019), https://doi.org/10.1007/978-3-030-26948-7_4
23. 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), https://doi.org/10.1007/978-3-031-58754-2_10
24. Qu, H., Xu, G.: On the provable dual attack for lwe by modulus switching. In: Hanaoka, G., Yang, B.Y. (eds.) *Advances in Cryptology – ASIACRYPT 2025*. pp. 34–64. Springer Nature Singapore, Singapore (2026)
25. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* **56**(6) (Sep 2009), <https://doi.org/10.1145/1568318.1568324>
26. Serre, J.P.: *A Course in Arithmetic*, vol. 7. Springer Science & Business Media (2012)
27. Stein, E.M., Shakarchi, R.: *Fourier analysis: An Introduction*, vol. 1. Princeton University Press (2003)
28. Tian, C., Liu, M., Xu, G.: Measure inequalities and the transference theorem in the geometry of numbers. *Proceedings of the American Mathematical Society* **142**(1), 47–57 (2014)
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, 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>
31. Yu, Y., Jia, H., Li, L., Ran, D., Qiu, Z., Zhang, S., Lin, X., Wang, X.: HuFu: Hash-and-Sign Signatures From Powerful Gadgets <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-1/spec-files/HuFu-spec-web.pdf>