Further Elaborations on the Provable Dual Attack

Kaveh Bashiri

Federal Office for Information Security, Germany,
name.lastname@bsi.bund.de

Abstract. The dual attack is one of the main strategies for attacking the learning-with-errors (LWE) problem. A recent line of work aims at provable analyses of dual attacks that rely on as few heuristic assumptions as possible.
We present a provable dual attack that recovers a subvector $\boldsymbol{s}_{\text{guess}}$ of the LWE secret, and we derive an explicit lower bound on its success probability without any heuristic assumption. Our analysis unifies the provable approaches of [35] and [9], and removes the very few informal assumptions left in the previous works. Inspired by [36], the dual vectors are sampled from the sublattice spanned by the first vectors of a reduced basis. Our key ingredient is a first-moment estimate for the wrong candidates, which is based on the exact uniform distribution of some appearing cosets. This leads us to discard certain typical assumptions on shortest vectors from previous works.
We apply our results to ML-KEM, under one heuristic on its structured matrix, and improve on the first provable dual attack by 22 to 52 bits without any assumption on the rank of the sublattice, and by 47 to 100 bits (21 to 56 bits compared with [36]) if this rank is bounded as in [36]. We also apply our results to ML-DSA, and demonstrate the attack experimentally on small instances. Our results do not threaten the security of ML-KEM or ML-DSA.

1 Introduction

The security of many important post-quantum schemes, such as ML-KEM [33] and ML-DSA [32], relies on the hardness of the *learning-with-errors (LWE)* problem or of one of its structured variants. Besides algebraic [6], combinatorial [24] and *primal lattice attacks* [18,23], a main strategy against LWE is the *dual attack*, on which we focus in this work. It successively guesses candidates $\tilde{\boldsymbol{s}}_{\text{guess}}$ for a subvector $\boldsymbol{s}_{\text{guess}}$ of the secret and evaluates each of them by a *score function*, which is built from short vectors of a dual lattice. With high probability, the score exceeds a certain *threshold* if and only if $\tilde{\boldsymbol{s}}_{\text{guess}} = \boldsymbol{s}_{\text{guess}}$.

Related work. The main idea of the dual attack originated in coding theory, and the first appearance in lattice theory seems to be [1]. Recent results

are [26,20,16,17,35,9,37,15,39,36]. For a long time, its analysis relied on the independence of the summands of the score function. In [16], it was shown that assumptions of this kind are inconsistent with practical observations, and a refined prediction of the score distribution was given in [17]. This triggered a series of works aiming at analyses that avoid such assumptions: [35] provide a first provable analysis for LWE, [9] a provable analysis fitting rather to the bounded distance decoding problem, [37] include the modulus switching technique into [35], and [39] revisit the CRT-based modulus switching of [37]. Most recently, [36] proposed to sample the dual vectors from the sublattice spanned by the first vectors of a reduced basis. Concurrently, [29] gave a provable analysis of the FFT-based dual attack over a general dual lattice with about n samples. A further line of work exploits the algebraic structure of schemes such as ML-KEM; e.g., [21] improve the primal hybrid attack and the code-based dual hybrid attack of [15].

Main contribution. We provide a *rigorous analysis of the dual attack* against LWE that relies on *no heuristic assumption*. In particular, we sample the dual vectors exactly from the discrete Gaussian distribution, by Klein's algorithm [25] followed by a rejection step (Lemma 5). This is even slightly cheaper than the MCMC sampler of [40] used in [35,39,36]. Our main results are explicit bounds on the success probability (Theorem 7) and on the complexity (Theorem 9), subject to an explicit inequality in the parameters of the attack ((11) for $t = m$), which does not depend on the reduced basis. We also demonstrate the attack on small instances (Section 7).

Main idea. Following [36], we sample the dual vectors from the sublattice Λ_t spanned by the first t vectors of a reduced basis of the dual lattice. The main challenge in the analysis of the attack is to control the wrong candidates $\tilde{\boldsymbol{s}}_{\mathrm{guess}} \neq \boldsymbol{s}_{\mathrm{guess}}$. In the previous works, an integral step here is to bound the distance of the target vector of such a candidate from the lattice from below by the length of the shortest nonzero vector of a related lattice minus the length of the error (cf. the proof of [35, Lemma 9]). This length is then estimated by the Gaussian heuristic. We, instead, use that the relevant coset of a wrong candidate is *exactly* uniformly distributed (Lemma 18), and that this uniformity is preserved under the projection. This enables us to control all wrong candidates simultaneously with standard tools from probability theory; see Lemma 17 and Lemma 20. Neither a tail bound for the discrete Gaussian nor a lower bound on a distance is needed, and of the reduced basis only $\det(\Lambda_t)$ enters. The attacker therefore chooses t as the smallest rank for which (11) holds, which only requires the Gram–Schmidt norms of the reduced basis. Our bound on the success probability holds for every reduced basis. The quality of the reduced basis affects only the cost, and already for $t = m$ our parameter constraint is (under the Gaussian heuristic) weaker than the one of [35,37,39]; cf. Remark 8.

Unification and rigor. We transfer the conditional-expectation framework of [9] to the LWE setting of [35]. Conditioning on the LWE sample separates its randomness from that of the sampled dual vectors, which yields *explicit* bounds

on the success probability in Theorem 7 and (13) (cf. the unspecified constant in the exponent in [39, Lemma 3]). Moreover, we remove the very few informal assumptions left in [35,37,39]. The assumption that the error does not exceed its mean (see e.g. [37, equation (12)]) is replaced by a concentration estimate (Lemma 16). The requirement that the matrix of the LWE sample has significantly more rows than columns, which fails for ML-KEM, is circumvented, as suggested in [35, Section 7], by sampling from a related q-ary lattice of the required shape (Section 3.1). Finally, as the lattice from which the dual vectors are drawn depends on the random matrix of the LWE sample, we make explicit that sampling from it is meaningful only conditionally on this matrix (Remark 4).

Guessing strategy. We incorporate the *improved guessing strategy* of [5,8], recently exploited in [23], which enumerates the candidates in decreasing order of probability. Using this technique, our attack guesses many more coordinates, and the dimension of the lattice decreases accordingly; cf. Remark 11. In Section 4 we add another layer to this strategy and obtain an improvement to our attack: one part of the guessed coordinates is enumerated in decreasing order of probability, the other is treated by the discrete Fourier transform.

Application to ML-KEM and ML-DSA. After invoking a heuristic (Heuristic 1), which accounts for the algebraic module structure used in ML-KEM, we apply the attack of Section 4 to its three parameter sets (Section 5). Without any assumption on the rank of the sublattice, that is, for $t = m$, it improves upon the provable dual attacks of [35,37,39] by 21.6, 36.2 and 52.2 bits (Table 1). If we bound this rank by assuming a generalized form of the *geometric series assumption* (Heuristic 2), it is 21.2, 42.2 and 56.2 bits cheaper than reported in [36], and 47.3, 72.1 and 99.9 bits cheaper than [35] (Table 2). The remaining differences to [36] are discussed in Remark 12. In all cases, our bound on the success probability holds without any assumption on the shape of the reduced basis. With minimal adaptations, the attack also applies to ML-DSA (Section 6).

Future research. The most immediate open problem is to discard Heuristic 1. A further one is to bound the determinant of the sublattice rigorously, which would allow us to discard Heuristic 2. Whether modulus switching can be made profitable within a provable framework is currently open; cf. [39, Table 1], where the provable attacks with modulus switching [37,39] are more expensive than [35]. Another goal is to replace our sampler by a practical lattice sieve: its output is confined to a sublattice of lower dimension, which our analysis already covers, but its distribution remains to be modeled rigorously; cf. [35, Section 7]. Finally, our analysis should be applied to more LWE-based schemes.

Outlook. After the preliminaries in Section 2, we describe the attack and state our main results in Section 3, and present the improved attack in Section 4. We apply it to ML-KEM in Section 5 and to ML-DSA in Section 6, and demonstrate the attack experimentally in Section 7. The proofs are given in Section 8 and Section 9.

2 Preliminaries

This preparatory section is organized as follows. We introduce some notation in Section 2.1 and provide a glimpse into lattice theory in Subsection 2.2. Then we introduce the widely used concept in probability theory of conditional expectation in Section 2.3. We introduce the LWE-problem in Subsection 2.4.

2.1 Notation

Throughout this paper $q \in \mathbb{N}$ denotes a prime integer, and $m, n \in \mathbb{N}$ satisfy $m \geq n$. We write vectors and matrices in bold letters. We write $\boldsymbol{B}^\mathsf{T}$ for the transpose of a matrix $\boldsymbol{B} \in \mathbb{R}^{m \times n}$, and denote by $\boldsymbol{I}_n$ the identity matrix in $\mathbb{R}^{n \times n}$. For $\boldsymbol{t}, \boldsymbol{t}' \in \mathbb{R}^m$, we denote by $\langle \boldsymbol{t}, \boldsymbol{t}' \rangle$ the Euclidean scalar product, by $\|\boldsymbol{t}\|$ the Euclidean norm and by $\mathrm{dist}(\boldsymbol{t}, L) = \inf_{\boldsymbol{v} \in L} \|\boldsymbol{t} - \boldsymbol{v}\|$ the distance of $\boldsymbol{t}$ from a measurable set $L \subset \mathbb{R}^m$. We write $\mathbb{Z}_q = \{0, 1, \ldots, q-1\}$ for the ring of integers modulo q. For $\eta \in \mathbb{N}$, we write $[\pm\eta]$ for the set $\{-\eta, -\eta+1, \ldots, \eta-1, \eta\}$. We write log and ln for the base-2 logarithm and the natural logarithm, respectively. For a measurable subset L, the *indicator function* $\mathbb{1}_L$ is defined by $\mathbb{1}_L(\boldsymbol{x}) = 1$ if $\boldsymbol{x} \in L$ and $\mathbb{1}_L(\boldsymbol{x}) = 0$ otherwise. Moreover, we use the classic and well-known notion of *Rényi entropy* in this paper: For a probability distribution $\mathcal{D}$ on some space L and for $\alpha > 0$ with $\alpha \neq 1$, the Rényi entropy of $\mathcal{D}$ is defined as

$$\mathcal{H}_\alpha(\mathcal{D}) := \tfrac{1}{1-\alpha} \log \left(\textstyle\sum_{\boldsymbol{x} \in L} p(\boldsymbol{x})^\alpha \right),$$

where, here, we denoted by p the probability mass function of $\mathcal{D}$. Note that $\mathcal{H}_\alpha(\mathcal{D}^{\otimes n}) = n\mathcal{H}_\alpha(\mathcal{D})$ for product distributions $\mathcal{D}^{\otimes n}$.

2.2 Lattices

In the following we only provide a rough glimpse into lattice theory. For more details we refer to [34], [30], [12] or [31].

A *lattice* $\mathcal{L} \subset \mathbb{R}^m$ is a set of the form $\mathcal{L} = \mathbb{Z}\boldsymbol{b}_1 + \cdots + \mathbb{Z}\boldsymbol{b}_t$, where $t \in \{1, \ldots, m\}$ and $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_t \in \mathbb{R}^m$ are linearly independent. We call $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_t$ a *basis* of $\mathcal{L}$ and t the *rank* of $\mathcal{L}$, and we say that $\mathcal{L}$ has *full rank* if $t = m$. If $\boldsymbol{B} \in \mathbb{R}^{m \times t}$ has $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_t$ as columns, the *determinant* (or *volume*) of $\mathcal{L}$ is defined by $\det(\mathcal{L}) := \sqrt{\det(\boldsymbol{B}^\mathsf{T}\boldsymbol{B})}$, which does not depend on the choice of the basis. The *dual lattice* of $\mathcal{L}$ is defined by $\widehat{\mathcal{L}} := \{\boldsymbol{w} \in \mathrm{span}_{\mathbb{R}}(\mathcal{L}) \mid \langle \boldsymbol{w}, \mathcal{L} \rangle \subset \mathbb{Z}\}$. It is a lattice of the same rank as $\mathcal{L}$, with $\det(\widehat{\mathcal{L}}) = \det(\mathcal{L})^{-1}$. Note that $\mathcal{L}$ is the dual lattice of $\widehat{\mathcal{L}}$.

***q*-ary lattices.** A lattice $\mathcal{L}$ is an m-dimensional *q-ary lattice* if $q\mathbb{Z}^m \subset \mathcal{L} \subset \mathbb{Z}^m$. Let $\boldsymbol{A} \in \mathbb{Z}^{m \times n}$. In this paper, we study q-ary lattices of the following form:

$$L_q(\boldsymbol{A}) = \{\boldsymbol{x} \in \mathbb{Z}^m \;\mid\; \exists \boldsymbol{s} \in \mathbb{Z}^n \;:\; \boldsymbol{A}\boldsymbol{s} = \boldsymbol{x} \pmod q\}$$
$$L_q^\perp(\boldsymbol{A}) = \{\boldsymbol{x} \in \mathbb{Z}^m \;\mid\; \boldsymbol{A}^\mathsf{T}\boldsymbol{x} = 0 \pmod q\}$$

It is well-known that the following duality relations exist between these lattices: $\widehat{L_q(\boldsymbol{A})} = \frac{1}{q} L_q^\perp(\boldsymbol{A})$ and $\widehat{L_q^\perp(\boldsymbol{A})} = \frac{1}{q} L_q(\boldsymbol{A})$.

Discrete Gaussian distribution. The *Gaussian density function* $\rho_s : \mathbb{R}^m \to (0, \infty)$ is, for $s > 0$, defined by $\rho_s(\boldsymbol{x}) = \mathrm{e}^{-\pi \frac{\|\boldsymbol{x}\|^2}{s^2}}$. For a discrete subset $L \subset \mathbb{R}^m$ we write $\rho_s(L) := \sum_{\boldsymbol{x} \in L} \rho_s(\boldsymbol{x})$.

The *Gaussian distribution* $D_{\mathcal{L},s}$ *over* $\mathcal{L}$ *of width* s is defined by the following probability mass function $D_{\mathcal{L},s}(\boldsymbol{v}) = \frac{\rho_s(\boldsymbol{v})}{\rho_s(\mathcal{L})}$ for all $\boldsymbol{v} \in \mathcal{L}$.

The *periodic Gaussian function* $f_{\mathcal{L},s} : \mathbb{R}^m \to \mathbb{R}$ is defined by $f_{\mathcal{L},s}(\boldsymbol{t}) = \frac{\rho_s(\mathcal{L}+\boldsymbol{t})}{\rho_s(\mathcal{L})}$ for all $\boldsymbol{t} \in \mathbb{R}^m$. Note that $f_{\mathcal{L},s}(\boldsymbol{t}) = f_{\mathcal{L},s}(-\boldsymbol{t})$ for all $\boldsymbol{t} \in \mathbb{R}^m$ (as $\mathcal{L} = -\mathcal{L}$).

Poisson summation formula. A very useful tool to switch from a lattice to its dual lattice is given by the *Poisson summation formula*, which we state here only for the Gaussian density function.

Lemma 1. *Let* $\mathcal{L} \subset \mathbb{R}^m$ *be a lattice of rank* $t \in \{1, \ldots, m\}$*, let* $V := \mathrm{span}_{\mathbb{R}}(\mathcal{L})$*,* $\boldsymbol{t} \in V$ *and* $s > 0$*. Then,*

$$\sum_{\boldsymbol{v} \in \mathcal{L}} \rho_s(\boldsymbol{v} - \boldsymbol{t}) = \frac{s^t}{\det(\mathcal{L})} \sum_{\boldsymbol{w} \in \hat{\mathcal{L}}} \rho_{1/s}(\boldsymbol{w}) \mathrm{e}^{-2\pi i \langle \boldsymbol{t}, \boldsymbol{w} \rangle}.$$

Proof. This is [38, Chapter VII, Proposition 15], applied to $\mathcal{L}$, the real vector space V and to the function $\boldsymbol{x} \mapsto \rho_s(\boldsymbol{x} - \boldsymbol{t})$ on V. Here we used that the volume of $V/\mathcal{L}$ equals $\det(\mathcal{L})$, that $\hat{\mathcal{L}}$ lies in V as well, and that the Fourier transform of this function on V equals $\boldsymbol{w} \mapsto s^t \rho_{1/s}(\boldsymbol{w}) \mathrm{e}^{-2\pi i \langle \boldsymbol{t}, \boldsymbol{w} \rangle}$ as V has dimension t.

Projected lattices. The following fact, which is [36, Lemma 1] in our notation, relates the dual of a sublattice to a projection of the primal lattice.

Lemma 2. *Let* $\mathcal{L} \subset \mathbb{R}^m$ *be a full-rank lattice, let* $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_m$ *be a basis of* $\widehat{\mathcal{L}}$*, and let* $t \in \{1, \ldots, m\}$*. Set* $\mathcal{L}' := \mathbb{Z}\boldsymbol{c}_1 + \cdots + \mathbb{Z}\boldsymbol{c}_t$ *and* $V := \mathrm{span}_{\mathbb{R}}(\mathcal{L}')$*, and let* $\pi : \mathbb{R}^m \to V$ *be the orthogonal projection onto* V*. Then* $\widehat{(\mathcal{L}')} = \pi(\mathcal{L})$*. In particular,* $\pi(\mathcal{L})$ *is a lattice of rank* t *with* $\widehat{\pi(\mathcal{L})} = \mathcal{L}'$ *and* $\det(\pi(\mathcal{L})) = 1/\det(\mathcal{L}')$*.*

Proof. The case $t = m$ is trivial. For $t < m$ the first claim is [36, Lemma 1]. The remaining claims follow by using standard properties of the dual lattice.

2.3 Conditional expectation

An essential tool for our proofs is the concept of *conditional expectation*. For background on this topic, we refer to any standard textbook in probability theory, such as [11, Sections 33 and 34]. The following intuitive description is a useful first step into this topic: Given a random variable $\boldsymbol{\Gamma}$, the conditional expectation given $\boldsymbol{\Gamma}$, denoted by $\mathbb{E}[\cdot \mid \boldsymbol{\Gamma}]$, is the best prediction under the condition that $\boldsymbol{\Gamma}$ is known. In particular, it is itself a random variable. The corresponding *conditional probability measure* is, like the usual probability measure, constructed as the conditional expectation of the indicator function $\mathbb{1}_L$ of an event L. That is, $\mathbb{P}[L \mid \boldsymbol{\Gamma}] = \mathbb{E}[\mathbb{1}_L \mid \boldsymbol{\Gamma}]$.

While this topic contains a vast theory with many applications, we are interested only in the case where $\boldsymbol{\Gamma}$ is a finite random variable. In this case, $\mathbb{E}[\,\cdot \mid \boldsymbol{\Gamma}]$ can be described very easily: Assume that L is the state space of $\boldsymbol{\Gamma}$, and that there exist $J \in \mathbb{N}$ and $l_1, \ldots, l_J \in L$ such that $\{l_1, \ldots, l_J\}$ are pairwise distinct, $\mathbb{P}(\boldsymbol{\Gamma} = l_i) > 0$ for all i, and $\sum_{i=1}^{J} \mathbb{P}(\boldsymbol{\Gamma} = l_i) = 1$. Then, for any random variable $\boldsymbol{X}$,

$$\mathbb{E}[\boldsymbol{X} \mid \boldsymbol{\Gamma}](\cdot) = \sum_{i=1}^{J} \mathbb{E}[\boldsymbol{X} \mid \boldsymbol{\Gamma} = l_i]\, \mathbb{1}_{\boldsymbol{\Gamma} = l_i}(\cdot), \tag{1}$$

where $\mathbb{E}[\boldsymbol{X} \mid \boldsymbol{\Gamma} = l_i] = \mathbb{E}[\boldsymbol{X}\, \mathbb{1}_{\boldsymbol{\Gamma} = l_i}]\, \mathbb{P}[\boldsymbol{\Gamma} = l_i]^{-1}$ is the usual conditional expectation given the *event* $\{\boldsymbol{\Gamma} = l_i\}$. In other words, $\mathbb{E}[\boldsymbol{X} \mid \boldsymbol{\Gamma}]$ is equal to one of the quantities $\{\mathbb{E}[\boldsymbol{X} \mid \boldsymbol{\Gamma} = l_i]\}_i$ depending on which realization is taken by $\boldsymbol{\Gamma}$.

Some important facts on conditional expectation that we will use here are recalled in [9]. Especially important is that the conditional expectation preserves many essential properties of the usual expectation such as linearity and monotonicity; see [11, Sections 33 and 34]. Another important property is shown in Section B, which is a generalization of the fact that when $\boldsymbol{X}$ is independent of $\boldsymbol{\Gamma}$, the following identity holds for any sufficiently integrable function G:

$$\mathbb{E}[G(\boldsymbol{X}, \boldsymbol{\Gamma}) \mid \boldsymbol{\Gamma}](\cdot) = \mathbb{E}_{\boldsymbol{X}}[G(\boldsymbol{X}, \boldsymbol{\Gamma}(\cdot))]. \tag{2}$$

Here, the left-hand side represents a random variable (as required by the conditional expectation definition), while the right-hand side computes the expectation over $\boldsymbol{X}$ while holding $\boldsymbol{\Gamma}(\cdot)$ fixed as an argument to G.

In general, (1) and (2) hold only *almost surely*, as $\mathbb{E}[\boldsymbol{X} \mid \boldsymbol{\Gamma}]$ is determined up to sets of probability zero. In our setting, however, every random variable on which we condition takes values in a finite set, so that the right-hand side of (1) is defined everywhere and all identities we use hold pointwise. We therefore omit the qualifier *almost surely* throughout.

2.4 The learning-with-errors (LWE) problem

Definition 3. *Let $\bar{m} \in \mathbb{N}$ with $\bar{m} \geq n$, and let $\bar{\boldsymbol{A}}$ be uniformly distributed on $\mathbb{Z}_q^{\bar{m} \times n}$. That is, the $\bar{m}n$ entries of $\bar{\boldsymbol{A}}$ are independent and each of them is uniformly distributed on $\mathbb{Z}_q$. Let $\eta \in \mathbb{N}$ be such that $2\eta < q$. Let $\{e_i\}_{i=1}^{\bar{m}} \cup \{s_i\}_{i=1}^{n}$ be a family of independent random variables, identically distributed according to the* centered binomial distribution with parameter η*, i.e. $e_i + \eta$ and $s_j + \eta$ are binomially distributed with parameters 2η and $\frac{1}{2}$ for all $(i, j) \in \{1, \ldots, \bar{m}\} \times \{1, \ldots, n\}$. We denote this centered binomial distribution with parameter η by $\mathcal{C}(\eta)$, and we write $\boldsymbol{e} = (e_i)_{i=1}^{\bar{m}}$ and $\boldsymbol{s} = (s_i)_{i=1}^{n}$, which are consequently random vectors in $[\pm\eta]^{\bar{m}}$ and $[\pm\eta]^n$, respectively. Moreover, we suppose that $\bar{\boldsymbol{A}}$ is independent of $(\boldsymbol{s}, \boldsymbol{e})$.*

We then call the pair $(\bar{\boldsymbol{A}}, \boldsymbol{b} = \bar{\boldsymbol{A}}\boldsymbol{s} + \boldsymbol{e} \bmod q) \in \mathbb{Z}_q^{\bar{m} \times n} \times \mathbb{Z}_q^{\bar{m}}$ an LWE-sample*, and we denote the resulting distribution by* $\mathrm{LWE}(\bar{m}, n, q, \eta)$*. In the* LWE-search-problem*, we are given an LWE-sample and asked to recover $\boldsymbol{s}$ (or a fraction thereof).*

3 The dual attack

We first describe the attack in Section 3.1. In Section 3.2, we describe the sampler of the dual vectors and the distribution of its output. We then formulate the main results of this paper: The estimate on the success probability (Section 3.3) and on the complexity (Section 3.4 and Section 3.5).

3.1 Description of the attack

Our attack is described in Algorithm 1. In the following we describe its strategy and provide some insights into the technical tools we use in the attack.

The strategy. Let $(\bar{\boldsymbol{A}}, \boldsymbol{b}) \in \mathbb{Z}_q^{\bar{m} \times n} \times \mathbb{Z}_q^{\bar{m}}$ be an LWE-sample, i.e. $\boldsymbol{b} = \bar{\boldsymbol{A}}\boldsymbol{s} + \boldsymbol{e} \bmod q$, where the random objects $\bar{\boldsymbol{A}}, \boldsymbol{s}$ and $\boldsymbol{e}$ are introduced in Definition 3.

Let $n = n_{\text{guess}} + n_{\text{dual}}$ for some $n_{\text{guess}}, n_{\text{dual}} \in \mathbb{N}$. We split $\bar{\boldsymbol{A}} \in \mathbb{Z}_q^{\bar{m} \times n}$ and $\boldsymbol{s} \in [\pm\eta]^n$, respectively, into the following two parts

$$\bar{\boldsymbol{A}} = \left(\bar{\boldsymbol{A}}_{\text{guess}}\ \bar{\boldsymbol{A}}_{\text{dual}}\right), \qquad \boldsymbol{s} = \begin{pmatrix} \boldsymbol{s}_{\text{guess}} \\ \boldsymbol{s}_{\text{dual}} \end{pmatrix},$$

where $\bar{\boldsymbol{A}}_{\text{dual}} \in \mathbb{Z}_q^{\bar{m} \times n_{\text{dual}}}, \bar{\boldsymbol{A}}_{\text{guess}} \in \mathbb{Z}_q^{\bar{m} \times n_{\text{guess}}}, \boldsymbol{s}_{\text{dual}} \in [\pm\eta]^{n_{\text{dual}}}, \boldsymbol{s}_{\text{guess}} \in [\pm\eta]^{n_{\text{guess}}}$. In this paper we aim to recover $\boldsymbol{s}_{\text{guess}}$.

The rough idea of our attack Algorithm 1 is to first successively guess candidates $\tilde{\boldsymbol{s}}_{\text{guess}} \in [\pm\eta]^{n_{\text{guess}}}$. Then, we evaluate this candidate by using it as an input to a so-called *score function*. If the output value of this score function is lower or equal to a certain *threshold* $c \cdot \mathcal{T} > 0$, we know with high probability that $\tilde{\boldsymbol{s}}_{\text{guess}} \neq \boldsymbol{s}_{\text{guess}}$. However, if the output value is greater than $c \cdot \mathcal{T}$, then, with high probability, we deduce that $\tilde{\boldsymbol{s}}_{\text{guess}} = \boldsymbol{s}_{\text{guess}}$, i.e. we recovered $\boldsymbol{s}_{\text{guess}}$.

The sampled dual vectors. It is obvious that the fundamental tool of our attack is the score function, so that the bulk part of our attack analysis will be dedicated to investigating its statistical properties. This score function, in the following denoted by g_W, depends on the sampling of a family $\{\boldsymbol{w}_1, \ldots, \boldsymbol{w}_N\} =: W$ of vectors *from the dual of a suitable lattice*. In this sampling process we modify some details with respect to the earlier or concurrent work such as [26,16,17,35,9,37,21,39,36]. These modifications have some advantages. In particular, these will make a formal application of our attack to ML-KEM ([33]) more rigorous.

Specifically, while in the papers [35,37,39] the dual vectors are sampled from $L_q^{\perp}(\bar{\boldsymbol{A}}_{\text{dual}})$, our sampling procedure is based on the lattice

$$L_q^{\perp}(\boldsymbol{A}_{\text{dual}}), \qquad \text{where } \boldsymbol{A}_{\text{dual}} := \begin{pmatrix} \bar{\boldsymbol{A}}_{\text{dual}} \\ \boldsymbol{I}_{n_{\text{dual}}} \end{pmatrix} \in \mathbb{Z}_q^{(\bar{m}+n_{\text{dual}}) \times n_{\text{dual}}} \tag{3}$$

The reason behind this (which is already mentioned in [35, Section 7]) is that the matrix $\bar{\boldsymbol{A}}$ is a square matrix in the case of ML-KEM. Consequently, $\bar{\boldsymbol{A}}_{\text{dual}}$ has

only slightly more rows than columns, since, in practice, n_{dual} is not much smaller than n. However, this is in contradiction to some parts of the attack analysis in [35,37,39], where the matrix is needed to have significantly more rows than columns. As suggested in [35, Section 7], switching to $L_q^\perp(\boldsymbol{A}_{\text{dual}})$ resolves this technical issue as $\boldsymbol{A}_{\text{dual}}$ has $\bar{m}$ more rows than columns. It is, on the other hand, important to note that this change requires a review of the analysis. This is one of the contributions of this paper. To sum up, the matrix, which generates the q-ary lattice in our paper has

$$m := \bar{m} + n_{\text{dual}} \tag{4}$$

rows and n_{dual} columns.

Note that the same lattice is used in [36]. Their parameters for ML-KEM, however, use 690, 810 and 1180 LWE samples, whereas ours use the $\bar{m} = n \in \{512, 768, 1024\}$ samples that ML-KEM provides; cf. Remark 12. However, we adapted another technical detail from [36], which we explain in the following.

We do not sample from the whole lattice $L_q^\perp(\boldsymbol{A}_{\text{dual}})$, but, following [36], from a sublattice of it. Let $\boldsymbol{B} = (\boldsymbol{b}_1, \dots, \boldsymbol{b}_m)$ be a basis of $L_q^\perp(\boldsymbol{A}_{\text{dual}})$, and let $\widetilde{\boldsymbol{b}}_1, \dots, \widetilde{\boldsymbol{b}}_m$ denote its Gram–Schmidt vectors. We assume that $\boldsymbol{B}$ is computed from $\bar{\boldsymbol{A}}_{\text{dual}}$ by a deterministic algorithm, so that $\boldsymbol{B}$ is a function of $\bar{\boldsymbol{A}}_{\text{dual}}$. For $t \in \{1, \dots, m\}$ (not to be confused with the vectors $\boldsymbol{t}$ below) we set

$$\Lambda_t := \mathbb{Z}\boldsymbol{b}_1 + \cdots + \mathbb{Z}\boldsymbol{b}_t \ \subset \ L_q^\perp(\boldsymbol{A}_{\text{dual}}) \qquad \text{and} \qquad V := \operatorname{span}_{\mathbb{R}}(\Lambda_t), \tag{5}$$

and we write $\pi : \mathbb{R}^m \to V$ for the orthogonal projection onto V. Note that Λ_t, V and π are functions of $\bar{\boldsymbol{A}}_{\text{dual}}$, and that $\det(\Lambda_t) = \prod_{i=1}^t \|\widetilde{\boldsymbol{b}}_i\|$. Moreover, for $t = m$ we have $\Lambda_m = L_q^\perp(\boldsymbol{A}_{\text{dual}})$, $\pi = \mathrm{id}$ and $\det(\Lambda_m) = q^{n_{\text{dual}}}$, where the latter holds for every realization of $\bar{\boldsymbol{A}}$, as the lower block of $\boldsymbol{A}_{\text{dual}}$ is $\boldsymbol{I}_{n_{\text{dual}}}$. In our analysis the rank t is not a free parameter: it is chosen as a function of $\bar{\boldsymbol{A}}_{\text{dual}}$, namely as the smallest rank for which the inequality (11) holds; cf. Theorem 7. Then, applying Lemma 2 both with $\mathcal{L} = L_q(\boldsymbol{A}_{\text{dual}}), \mathcal{L}' = \frac{1}{q}\Lambda_t \subset \frac{1}{q}L_q^\perp(\boldsymbol{A}_{\text{dual}})$ and $\mathcal{L} = \mathbb{Z}^m, \mathcal{L}' = \mathbb{Z}^m \cap V$ yields that

$$\pi(L_q(\boldsymbol{A}_{\text{dual}})) = \widehat{(\tfrac{1}{q}\Lambda_t)}, \quad \det(\pi(L_q(\boldsymbol{A}_{\text{dual}}))) = \tfrac{q^t}{\det(\Lambda_t)}, \quad \widehat{\pi(\mathbb{Z}^m)} = \mathbb{Z}^m \cap V, \tag{6}$$

and that $\pi(L_q(\boldsymbol{A}_{\text{dual}}))$ and $\pi(\mathbb{Z}^m)$ are lattices of rank t. Here we have tacitly used that $\mathbb{Z}^m \cap V$ contains Λ_t and thus has rank t, and that, being the intersection of $\mathbb{Z}^m$ with a subspace, it is primitive in $\mathbb{Z}^m$ (i.e. a basis of it can be extended to a basis of $\mathbb{Z}^m$).

In conclusion, as a first step in the construction of the score function, we (independently) sample $N \in \mathbb{N}$ vectors $\{\boldsymbol{w}_1, \dots, \boldsymbol{w}_N\} =: W$ from the lattice Λ_t. We instantiate this sampling step by N independent runs of the exact sampler of Section 3.2 on the basis $(\boldsymbol{b}_1, \dots, \boldsymbol{b}_t)$ of Λ_t. The distribution of the output vectors is described in Lemma 5.

The score function. Then, a reasonable candidate for the score function, as motivated in previous works such as [26,16,17,35,9,37,21,39,36], is given by

$$g_W(\boldsymbol{t}) = \frac{1}{N}\sum_{j=1}^{N} \cos\left(2\pi \frac{1}{q}\langle \boldsymbol{t}, \boldsymbol{w_j}\rangle\right) \qquad \text{for } \boldsymbol{t} \in \mathbb{Z}^{\bar{m}+n_{\text{dual}}} = \mathbb{Z}^m \tag{7}$$

In our attack, we will, for each candidate $\tilde{\boldsymbol{s}}_{\text{guess}}$, apply this function g_W with

$$\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}}) := \begin{pmatrix} \boldsymbol{b} - \bar{\boldsymbol{A}}_{\text{guess}}\tilde{\boldsymbol{s}}_{\text{guess}} \\ 0 \end{pmatrix} \in \mathbb{Z}^m \tag{8}$$

as input. Note that, as $\boldsymbol{b} = \bar{\boldsymbol{A}}\boldsymbol{s} + \boldsymbol{e} = \bar{\boldsymbol{A}}_{\text{dual}}\boldsymbol{s}_{\text{dual}} + \bar{\boldsymbol{A}}_{\text{guess}}\boldsymbol{s}_{\text{guess}} + \boldsymbol{e} \bmod q$, we have that

$$\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}}) = \begin{pmatrix} \boldsymbol{e} + \bar{\boldsymbol{A}}_{\text{dual}}\boldsymbol{s}_{\text{dual}} + \bar{\boldsymbol{A}}_{\text{guess}}(\boldsymbol{s}_{\text{guess}} - \tilde{\boldsymbol{s}}_{\text{guess}}) \\ 0 \end{pmatrix} \bmod q \tag{9}$$

As a consequence, it is important to observe that $\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}})$ is fully determined by the random vector $(\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})$ and by the vector $\tilde{\boldsymbol{s}}_{\text{guess}}$. The latter, in our attack, will be deterministic (as we choose $\tilde{\boldsymbol{s}}_{\text{guess}}$ deterministically according to our guessing strategy, which we will explain in the following). We study the statistical properties of this score function evaluated at these vectors in the following sections.

The guessing strategy. Our guessing strategy also slightly deviates from [35,37,39]. They chose the candidate keys $\tilde{\boldsymbol{s}}_{\text{guess}}$ deterministically according to a not-further specified enumeration of the corresponding space, thereby obtaining a framework, which is suitable for all possible secret vector distributions. We, on the other hand, aim to exploit the specific non-uniform distribution of the secret in ML-KEM in this guessing step. This way we can benefit from some speed-up in the guessing complexity in comparison to guessing for uniform distributions; see [23], where this idea was already exploited.

Recall that $\mathcal{C}(\eta)$ denotes the centered binomial distribution with parameter η, and that $\boldsymbol{s}_{\text{guess}}$ is distributed according to the product distribution $\mathcal{C}(\eta)^{\otimes n_{\text{guess}}}$. Note that the state space of $\boldsymbol{s}_{\text{guess}}$ is given by $[\pm\eta]^{n_{\text{guess}}}$. Already since the early work by [5], it is known that the optimal guessing strategy is to successively check all possible realization values *in decreasing order of probability*, until the correct one is found; see [8] for more details.

This guessing strategy requires access to an algorithm GetRealization, which, on input $i \in \{1, \ldots, |[\pm\eta]^{n_{\text{guess}}}| = (2\eta+1)^{n_{\text{guess}}}\}$, efficiently outputs the i-th most likely realization for the random vector $\boldsymbol{s}_{\text{guess}}$. In other words, GetRealization outputs the vector $\boldsymbol{u_i} \in [\pm\eta]^{n_{\text{guess}}}$ such that the probability $\mathbb{P}[\boldsymbol{s}_{\text{guess}} = \boldsymbol{u_i}]$ is the i-th highest probability among all probabilities corresponding to elements of $[\pm\eta]^{n_{\text{guess}}}$. For our setting, where we consider a product distribution over some finite set, it is shown in [14] that an efficient instantiation of such an algorithm exists. The corresponding time complexity of this algorithm is, in our setting, given by $\mathcal{O}((2\eta+1)\cdot n_{\text{guess}})$ (after one initial pre-computation phase running in time $\widetilde{\mathcal{O}}((n_{\text{guess}})^{2\eta})$). Throughout the paper, we assume black box access to GetRealization.

The threshold parameter. The threshold parameter $c \cdot \mathcal{T}$ is essential for our attack as it provides a boundary such that, with high probability, $g_W(\boldsymbol{t}(\boldsymbol{s}_{\text{guess}}))$ falls on one side while $g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}}))$ for $\tilde{\boldsymbol{s}}_{\text{guess}} \neq \boldsymbol{s}_{\text{guess}}$ remains on the other. Specifically, we set

$$c \in (0,1) \quad \text{and} \quad \mathcal{T} = \mathrm{e}^{-t\pi s^2 \frac{\eta}{2}(1+\gamma_t)}, \text{ where } \gamma_t = 2\sqrt{\tfrac{\ln(2)}{t}} + \tfrac{2\ln(2)}{t}. \tag{10}$$

Here, $s > 0$ is the width parameter of the sampled dual vectors; cf. Lemma 5. Note that $\gamma_t \in O(1/\sqrt{t})$. Moreover, note that the exponent of $\mathcal{T}$ is proportional to the rank t of the sublattice and not to the dimension m. This is the benefit of sampling from Λ_t instead of $L_q^{\perp}(\boldsymbol{A}_{\text{dual}})$. $\mathcal{T}$ is derived from a precise analysis (in Section 8 and Section 9) of the conditional expectation of g_W (evaluated at both the correct secret $\boldsymbol{s}_{\text{guess}}$ and incorrect candidates $\tilde{\boldsymbol{s}}_{\text{guess}} \neq \boldsymbol{s}_{\text{guess}}$) given the LWE sample.

Algorithm 1 The dual attack

Input: An LWE-sample $(\bar{\boldsymbol{A}}, \boldsymbol{b})$ according to Definition 3.
Input: $t \in \{1, \ldots, m\}$ and $N \in \mathbb{N}$ vectors $\{\boldsymbol{w}_1, \ldots, \boldsymbol{w}_N\} =: W$ from the lattice Λ_t sampled as in Lemma 5.
Input: Threshold parameters $c, \mathcal{T} > 0$ defined in (10).
Output: Candidate $\tilde{\boldsymbol{s}}_{\text{guess}}$ for $\boldsymbol{s}_{\text{guess}}$ or $\perp$.
for $i \leftarrow 1$ **to** $(2\eta+1)^{n_{\text{guess}}}$ **do**
    $\tilde{\boldsymbol{s}}_{\text{guess}} \leftarrow \textsc{GetRealization}(i)$
    if $g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}})) > c \cdot \mathcal{T}$ **then** ▷ g_W and $\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}})$ are defined in (7) and (8).
        return $\tilde{\boldsymbol{s}}_{\text{guess}}$
    end if
end for
return $\perp$

3.2 The distribution of the dual vectors

As mentioned above, our attack relies on sampling from the lattice Λ_t, which we defined in (5). In this subsection, we describe the sampler and the distribution of its output. Before we do so, we address a point of rigor that arises as soon as $\bar{\boldsymbol{A}}$ is random, and which, to our knowledge, has not been made explicit in previous works on the dual attack.

Remark 4. Recall that in standardized KEMs such as ML-KEM [33], the matrix $\bar{\boldsymbol{A}}$ is sampled at random, and we, as in related works such as [35,37,39], incorporated randomness of $\bar{\boldsymbol{A}}$ into Definition 3. Now suppose we wish to sample $\boldsymbol{w}$ from $D_{\Lambda_t, qs}$. As $\bar{\boldsymbol{A}}$ is random, so is Λ_t. The definition of the vector $\boldsymbol{w}$ is therefore meaningful only with respect to $\mathbb{P}[\cdot \mid \bar{\boldsymbol{A}}]$, the conditional probability measure given $\bar{\boldsymbol{A}}$. Accordingly, a quantity such as $\mathbb{P}[\boldsymbol{w} \in L]$ for a measurable $L \subseteq \mathbb{R}^m$

has to be computed as

$$\mathbb{P}\left[\boldsymbol{w} \in L\right] = \mathbb{E}\left[\ \mathbb{P}\left[\boldsymbol{w} \in L \mid \bar{\boldsymbol{A}}\right]\ \right] = \mathbb{E}\left[\ \sum_{\boldsymbol{v}\in L\cap\Lambda_t} \frac{\rho_{qs}(\boldsymbol{v})}{\rho_{qs}(\Lambda_t)}\ \right],$$

where, on the right-hand side, the inner expression is a random variable, being a function of $\bar{\boldsymbol{A}}$, and the outer expectation is taken over the randomness of $\bar{\boldsymbol{A}}$. Every statement about objects sampled from Λ_t is to be understood in this sense, i.e. by invoking $\mathbb{P}[\cdot \mid \bar{\boldsymbol{A}}]$.

The sampler. Let $t \in \{1,\ldots,m\}$, let $s > 0$, and set $\sigma_i := qs/\|\widetilde{\boldsymbol{b}}_i\|$ for $i = 1,\ldots,t$. We use Klein's randomized nearest plane algorithm [25,19] on $(\boldsymbol{b}_1,\ldots,\boldsymbol{b}_t)$ with width qs as a proposal, and accept its output with probability $\prod_{i=1}^{t} \rho_{\sigma_i}(\mathbb{Z}-c_i)/\rho_{\sigma_i}(\mathbb{Z})$, where $c_1,\ldots,c_t$ are the centers used by Klein's algorithm. Otherwise, we repeat with fresh randomness. As in the proof of [13, Lemma 2.3], whose argument applies to every width, the accepted output is distributed according to $D_{\Lambda_t,qs}$, and one round is accepted with probability $\rho_{qs}(\Lambda_t)/\prod_{i=1}^{t}\rho_{\sigma_i}(\mathbb{Z}) = \Delta_t$, where Δ_t is defined in (15). In particular, the expected number of calls to Klein's algorithm per output vector equals $1/\Delta_t$.

Lemma 5. *Let $(\bar{\boldsymbol{A}}, \boldsymbol{b})$ be an LWE-sample according to Definition 3, let $s > 0$, and let $t \in \{1,\ldots,m\}$ and $N \in \mathbb{N}$ be functions of $\bar{\boldsymbol{A}}_{\mathrm{dual}}$. Let $\{\boldsymbol{w}_1,\ldots,\boldsymbol{w}_N\} =: W$ be the outputs of N independent runs of the above sampler on $(\boldsymbol{b}_1,\ldots,\boldsymbol{b}_t)$. Then:*

- *For all $i = 1,\ldots,N$, the conditional distribution of $\boldsymbol{w}_i$ given $\bar{\boldsymbol{A}}$ is $D_{\Lambda_t,qs}$. That is, for every measurable $L \subset \mathbb{R}^m$,*

$$\mathbb{P}\left[\boldsymbol{w}_i \in L \mid \bar{\boldsymbol{A}}\right] \ = \ \sum_{\boldsymbol{v}\in L\cap\Lambda_t} \frac{\rho_{qs}(\boldsymbol{v})}{\rho_{qs}(\Lambda_t)}.$$

- *Conditionally on $\bar{\boldsymbol{A}}$ (i.e. with respect to the probability measure $\mathbb{P}[\cdot \mid \bar{\boldsymbol{A}}]$), the family $\{\boldsymbol{s}\} \cup \{\boldsymbol{e}\} \cup \left(\cup_{i=1}^{N}\{\boldsymbol{w}_i\}\right)$ is a family of independent random variables.*

In other words, the second item means that

$$\mathbb{E}\left[f(\boldsymbol{s}) \cdot f'(\boldsymbol{e}) \cdot \prod_{i=1}^{N} f_i(\boldsymbol{w}_i) | \bar{\boldsymbol{A}}\right] = \mathbb{E}\left[f(\boldsymbol{s})|\bar{\boldsymbol{A}}\right] \cdot \mathbb{E}\left[f'(\boldsymbol{e})|\bar{\boldsymbol{A}}\right] \cdot \prod_{i=1}^{N} \mathbb{E}\left[f_i(\boldsymbol{w}_i)|\bar{\boldsymbol{A}}\right]$$

for all bounded measurable functions $f : [\pm\eta]^n \to \mathbb{R}$, $f' : [\pm\eta]^{\bar{m}} \to \mathbb{R}$ and $f_1,\ldots,f_N : \mathbb{R}^m \to \mathbb{R}$.

Proof. The input $(\boldsymbol{b}_1,\ldots,\boldsymbol{b}_t)$ of the sampler, as well as N, is a function of $\bar{\boldsymbol{A}}_{\mathrm{dual}}$, and the sampler uses randomness that is independent of $(\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})$ and across its runs. As $(\boldsymbol{s}, \boldsymbol{e})$ is independent of $\bar{\boldsymbol{A}}$, both claims follow from the properties of the sampler stated above.

The following lemma is a direct consequence of the above, and we will use this fact throughout our analysis.

Lemma 6. *Suppose the setting of Lemma 5, and let* $\tilde{\boldsymbol{s}}_{\text{guess}} \in [\pm\eta]^{n_{\text{guess}}}$. *Then the family*

$$\left\{\cos\left(2\pi\tfrac{1}{q}\langle \boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}}), \boldsymbol{w}_j\rangle\right)\right\}_{j=1,\dots,N}$$

is independent with respect to $\mathbb{P}[\,\cdot \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})]$.

Proof. By Lemma 5, the vector $(\boldsymbol{w}_1, \dots, \boldsymbol{w}_N)$ is independent of $(\boldsymbol{s}, \boldsymbol{e})$ with respect to $\mathbb{P}[\,\cdot \mid \bar{\boldsymbol{A}}]$. Moreover, $\cos(2\pi\frac{1}{q}\langle \boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}}), \boldsymbol{w}_j\rangle) = \cos(2\pi\frac{1}{q}\langle \boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}}) \bmod q, \boldsymbol{w}_j\rangle)$ and, by (9), we know that the vector $\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}}) \bmod q$ is a function of $(\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})$. Therefore, we are in the exact same setting of Corollary 24 (with $\boldsymbol{\Gamma} = \bar{\boldsymbol{A}}$, $\boldsymbol{\Theta} = (\boldsymbol{s}, \boldsymbol{e})$ and $\boldsymbol{X} = (\boldsymbol{w}_1, \dots, \boldsymbol{w}_N)$), which shows that the family $\{\cos(2\pi\frac{1}{q}\langle \boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}}), \boldsymbol{w}_j\rangle)\}_j$ is independent.

3.3 Success probability of Algorithm 1

Theorem 7. *Let* $(\bar{\boldsymbol{A}}, \boldsymbol{b})$ *be an LWE-sample according to Definition 3, and* $t \in \{1, \dots, m\}$. *Let* m *be defined as in* (4), Λ_t *as in* (5), *and* c, γ_t *and* $\mathcal{T}$ *as in* (10). *Consider the inequality*

$$\tfrac{2}{c\,\mathcal{T}} \cdot m\, 2^{n_{\text{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}\, \rho_s(\mathbb{Z})^m\, \det(\Lambda_t) \;\le\; (qs)^t, \tag{11}$$

where $\mathcal{H}_{1/2}$ *denotes the Rényi entropy with parameter* $1/2$ *(which we introduced in Section 2.1). Suppose that* (11) *holds for* $t = m$. *This is a deterministic condition on the parameters of the attack (as* $\det(\Lambda_m) = q^{n_{\text{dual}}}$, *which does not depend on* $\bar{\boldsymbol{A}}$*). Let* t^* *denote the smallest* $t \in \{1, \dots, m\}$ *for which* (11) *holds (recall that* $\det(\Lambda_t)$, *and hence* t^*, *is a function of* $\bar{\boldsymbol{A}}_{\text{dual}}$*). Moreover, let*

$$c \;=\; \tfrac{1}{2} \qquad \textit{and} \qquad N \;:=\; \left\lceil \tfrac{8}{\mathcal{T}^2}\Big(n_{\text{guess}}\,\mathcal{H}_{1/2}(\mathcal{C}(\eta))\ln(2) + \ln(m)\Big)\right\rceil, \tag{12}$$

where $\mathcal{T}$ *is evaluated at* $t = t^*$ *(so that* N *is a function of* $\bar{\boldsymbol{A}}_{\text{dual}}$*). In particular,* (11), *and hence* t^*, *is understood with* $c = \frac{1}{2}$. *Run Algorithm 1 with* $t = t^*$ *and with* N *dual vectors sampled as in Lemma 5. Then, Algorithm 1 is successful, (i.e. it outputs* $\boldsymbol{s}_{\text{guess}}$*) with probability at least* $\frac{1}{2} - \frac{3}{m}$.

Proof. Recall the definition of $\mathcal{C}(\eta)$ in Definition 3. For $i = 1, \dots, (2\eta+1)^{n_{\text{guess}}}$ let $\boldsymbol{u_i} := \textsc{GetRealization}(i) \in [\pm\eta]^{n_{\text{guess}}}$ and $p_i := \mathbb{P}[\boldsymbol{s}_{\text{guess}} = \boldsymbol{u_i}]$. That is, $p_1 \ge p_2 \ge \dots$, and p_i is the probability of the i-th likeliest realization of a random variable drawn from $\mathcal{C}(\eta)^{\otimes n_{\text{guess}}}$, where ties are handled as in GetRealization. Let $R : [\pm\eta]^{n_{\text{guess}}} \to \{i = 1, \dots, (2\eta+1)^{n_{\text{guess}}}\}$ be the corresponding *rank function* that outputs the position of a realization value with respect to this ordering. That is, for $\tilde{\boldsymbol{s}}_{\text{guess}} \in [\pm\eta]^{n_{\text{guess}}}$, $R(\tilde{\boldsymbol{s}}_{\text{guess}}) = i$ if and only if $\tilde{\boldsymbol{s}}_{\text{guess}} = \boldsymbol{u_i}$.

Then, Algorithm 1 fails to output the correct key $\boldsymbol{s}_{\text{guess}}$ only if one of the following events occurs:

(FN) *False negative:* the correct candidate is rejected, that is, $g_W(\boldsymbol{t}(\boldsymbol{s}_{\text{guess}})) \le c{\cdot}\mathcal{T}$.

(FP) *False positive:* a wrong candidate that is examined before the correct one is accepted, that is, there exists some $\tilde{\boldsymbol{s}}_{\text{guess}} \neq \boldsymbol{s}_{\text{guess}}$ such that

$$R(\tilde{\boldsymbol{s}}_{\text{guess}}) < R(\boldsymbol{s}_{\text{guess}}) \qquad \text{and} \qquad g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}})) > c \cdot \mathcal{T}$$

This implies that

$$\begin{aligned}
&\mathbb{P}[\text{Algorithm 1 does not output } \boldsymbol{s}_{\text{guess}}] \;\leq\; \mathbb{P}[\text{(FN) occurs}] + \mathbb{P}[\text{(FP) occurs}] \\
&= \mathbb{P}[g_W(\boldsymbol{t}(\boldsymbol{s}_{\text{guess}})) \leq c \cdot \mathcal{T}] \\
&\quad + \mathbb{P}[\,\exists \tilde{\boldsymbol{s}}_{\text{guess}} \in [\pm\eta]^{n_{\text{guess}}} : (R(\tilde{\boldsymbol{s}}_{\text{guess}}) < R(\boldsymbol{s}_{\text{guess}})) \wedge (g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}})) > c\mathcal{T})]
\end{aligned}$$

Applying Lemma 15 and Lemma 17 yields

$$\begin{aligned}
&\mathbb{P}\left[\text{Algorithm 1 does not output } \boldsymbol{s}_{\text{guess}}\right] \\
&\qquad \leq\; \mathbb{E}\big[\mathrm{e}^{-\frac{1}{2}N(1-c)^2\mathcal{T}^2}\big] + \tfrac{1}{2} + \tfrac{1}{m} + 2^{n_{\text{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}\,\mathbb{E}\big[\mathrm{e}^{-\frac{1}{2}Nc^2\mathcal{T}^2}\big],
\end{aligned} \tag{13}$$

where the expectations are taken with respect to the randomness of N and $\mathcal{T}$, which both depend on t^* and thus on $\bar{\boldsymbol{A}}_{\text{dual}}$ (recall that $\mathcal{T}$ is evaluated at $t = t^*$). With $c = \frac{1}{2}$ we have $\frac{1}{2}(1-c)^2 = \frac{1}{2}c^2 = \frac{1}{8}$, so the two exponential terms in (13) carry the same exponent $\frac{1}{8}N\mathcal{T}^2$. By (12),

$$\tfrac{1}{8}N\mathcal{T}^2 \;\geq\; n_{\text{guess}}\,\mathcal{H}_{1/2}(\mathcal{C}(\eta))\ln(2) + \ln(m),$$

so that $\mathrm{e}^{-\frac{1}{2}N(1-c)^2\mathcal{T}^2} \leq \frac{1}{m}$ and $2^{n_{\text{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}\,\mathrm{e}^{-\frac{1}{2}Nc^2\mathcal{T}^2} \leq \frac{1}{m}$ for every value of $\bar{\boldsymbol{A}}_{\text{dual}}$. Hence the right-hand side of (13) is at most $\frac{1}{2} + \frac{3}{m}$, which concludes the proof.

Remark 8. (i) **The role of the reduced basis.** The hypothesis of Theorem 7, that is, (11) for $t = m$, does not involve the reduced basis. The latter enters only through t^*. Hence, the quality of the reduced basis does not affect the success probability. The proofs in Section 8 and Section 9 only use that t^* is a function of $\bar{\boldsymbol{A}}_{\text{dual}}$ for which (11) holds for every realization of $\bar{\boldsymbol{A}}_{\text{dual}}$. Hence t^* may be replaced by any such t, for instance by $t = m$, and $t = t^*$ is the choice which maximises $\mathcal{T}$ and thus minimises N.

(ii) **The full lattice.** As $\Lambda_m = L_q^{\perp}(\boldsymbol{A}_{\text{dual}})$ and $\det(\Lambda_m) = q^{n_{\text{dual}}}$, for $t = m$, the hypothesis (11) of Theorem 7 does not depend on $\bar{\boldsymbol{A}}$, and (19) (hence (11)) holds for $t = m$ at the parameters of Table 1 and Table 2.

(iii) **The condition of [35,37,39].** The assumption in [35,37,39] corresponding to our assumption (11) says that the error lies inside the unique decoding radius. It is read with the typical length $\sqrt{m\eta/2}$ of $\left(\begin{smallmatrix}\boldsymbol{e}\\ -\boldsymbol{s}_{\text{dual}}\end{smallmatrix}\right)$ and with the Gaussian heuristic for the m-dimensional lattice $L_q\big((\boldsymbol{A}_{\text{guess}}\ \ \boldsymbol{A}_{\text{dual}})\big)$ of determinant $q^{\bar{m}-n_{\text{guess}}}$, where $\boldsymbol{A}_{\text{guess}} := \left(\begin{smallmatrix}\bar{\boldsymbol{A}}_{\text{guess}}\\ \boldsymbol{0}\end{smallmatrix}\right)$. We claim that, for $m \geq 9$, every instantiation which satisfies the corresponding assumption of [35,37,39] also satisfies (11) for $t = m$ and $s = (\pi\eta(1+\gamma_m))^{-1/2}$.
Indeed, the assumption of [35,37,39] then reads $q^{\bar{m}/m} > 2\sqrt{\pi e\eta}\,q^{n_{\text{guess}}/m}$, whereas, as $\mathcal{T} = \mathrm{e}^{-m/2}$ for this s and $c = \frac{1}{2}$, (11) for $t = m$ reads

$\sqrt{\pi e \eta(1+\gamma_m)}\,\rho_s(\mathbb{Z})\,(4m\,2^{n_{\text{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))})^{1/m} \le q^{\bar{m}/m}$. The former implies the latter, because $2^{\mathcal{H}_{1/2}(\mathcal{C}(\eta))} \le q$ and $\sqrt{1+\gamma_m}\,\rho_s(\mathbb{Z})\,(4m)^{1/m} \le 2$ for $m \ge 9$.

The converse fails, e.g. for $q = 3329$, $n = \bar{m} = 512$, $\eta = 3$ and $n_{\text{guess}} = 340$.

(iv) **The success probability.** The summand $\frac{1}{2}$ in (13) stems from Lemma 16. For $\delta \in (0,1)$, replacing $\ln(2)$ by $\ln(1/\delta)$ in the definition of γ_t yields a success probability of at least $1 - \delta - \frac{3}{m}$, for $\delta < \frac{1}{2}$ at the price of a larger N (and possibly a larger t^*). We chose $\delta = \frac{1}{2}$, as the parameters of [36] are chosen such that their success probability is approximately $\frac{1}{2}$; cf. [36, Section 4.3]. There, however, the projected error is replaced by its typical length (cf. Remark 12), whereas γ_t accounts for its deviation.

3.4 Complexity of Algorithm 1

Theorem 9. *Suppose the same setting as in Theorem 7, and that $m > 6$. Then the expected complexity of Algorithm 1, conditioned on the event that Algorithm 1 outputs $\boldsymbol{s}_{\text{guess}}$, is at most*

$$\frac{1}{\frac{1}{2}-\frac{3}{m}} \cdot \mathbb{E}\Big[T_{\text{sampling}}(N, qs) + \text{poly}(m,n) \cdot \min\left\{ N \cdot 2^{n_{\text{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))},\, N + q^{n_{\text{guess}}} \right\} \Big] + \widetilde{\mathcal{O}}\big(n_{\text{guess}}^{2\eta}\big),$$

where $T_{\text{sampling}}(N, qs)$ denotes the expected cost, given $\bar{\boldsymbol{A}}$, of sampling the N vectors $\{\boldsymbol{w}_1, \ldots, \boldsymbol{w}_N\}$ from Λ_{t^*}, and where the expectation is taken with respect to $\bar{\boldsymbol{A}}_{\text{dual}}$ (recall that t^* and N are functions of $\bar{\boldsymbol{A}}_{\text{dual}}$).

Proof. The pre-computation phase of GetRealization is carried out once, at cost $\widetilde{\mathcal{O}}(n_{\text{guess}}^{2\eta})$; cf. [14]. The N vectors are sampled once, and, given $\bar{\boldsymbol{A}}$, this has expected cost $T_{\text{sampling}}(N, qs)$, which is a function of $\bar{\boldsymbol{A}}_{\text{dual}}$.

The remaining term corresponds to the cost of the loop in Algorithm 1, where we find two ways to upper bound this term. Depending on the parameter set, either of them may be favorable. Of course, the true average complexity of the loop is lower or equal to the minimum of both.

For the first way, recall the definition of the rank function R from the proof of Theorem 7 and observe that, on the event S that Algorithm 1 outputs $\boldsymbol{s}_{\text{guess}}$, the loop performs exactly $R(\boldsymbol{s}_{\text{guess}})$ iterations. In each iteration, evaluating $g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{guess}}))$ as in (7) costs $\mathcal{O}(N\,m)$ arithmetic operations, and one call to GetRealization costs $\mathcal{O}((2\eta+1)n_{\text{guess}})$. Hence, on S, the loop costs at most $\text{poly}(m,n) \cdot N \cdot R(\boldsymbol{s}_{\text{guess}})$. As N is a function of $\bar{\boldsymbol{A}}_{\text{dual}}$, and as $\boldsymbol{s}_{\text{guess}}$ is independent of $\bar{\boldsymbol{A}}_{\text{dual}}$,

$$\mathbb{E}\left[N \cdot R(\boldsymbol{s}_{\text{guess}}) \;\middle|\; \bar{\boldsymbol{A}}_{\text{dual}} \right] \;=\; N \cdot \mathbb{E}\left[R(\boldsymbol{s}_{\text{guess}}) \right] \;=\; N \cdot \textstyle\sum_k k \cdot p_k$$

where we introduced the probabilities $(p_k)_k$ in the beginning of the proof of Theorem 7. Using Arikan's classical inequalities on guessing (see [8, Theorem 3.2] or [5]), yields that $\mathbb{E}\left[R(\boldsymbol{s}_{\text{guess}}) \right] \le 2^{n_{\text{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}$.

For the second way to estimate the loop cost, we, as in [35], use the DFT trick of [20]: all scores $g_W(\boldsymbol{t}(\boldsymbol{x}))$, $\boldsymbol{x} \in \mathbb{Z}_q^{n_{\text{guess}}}$, are computed by one discrete Fourier transform, after which the loop of Algorithm 1 is carried out with these precomputed scores. Hence the loop costs at most $\text{poly}(m,n) \cdot (N + q^{n_{\text{guess}}})$ in this variant.

To sum up, let the loop be carried out in the variant for which the respective bound, $N \cdot 2^{n_{\text{guess}} \mathcal{H}_{1/2}(\mathcal{C}(\eta))}$ or $N + q^{n_{\text{guess}}}$, is smaller (both are functions of $\bar{\boldsymbol{A}}_{\text{dual}}$). Both variants examine the same candidates in the same order and with the same scores, so that the output of Algorithm 1, and hence S, does not depend on this choice. Then, on S and up to the term $\widetilde{\mathcal{O}}(n_{\text{guess}}^{2\eta})$, the complexity of Algorithm 1 is at most a random variable $X \geq 0$ with $\mathbb{E}[X \mid \bar{\boldsymbol{A}}_{\text{dual}}] \leq T_{\text{sampling}}(N, qs) + \text{poly}(m,n) \cdot \min\{N \cdot 2^{n_{\text{guess}} \mathcal{H}_{1/2}(\mathcal{C}(\eta))}, N + q^{n_{\text{guess}}}\}$. As $\mathbb{P}[S] \geq \frac{1}{2} - \frac{3}{m}$ by Theorem 7, we conclude with $\mathbb{E}[X \mid S] = \frac{\mathbb{E}[\mathbb{1}_S X]}{\mathbb{P}[S]} \leq \left(\frac{1}{2} - \frac{3}{m}\right)^{-1} \mathbb{E}[X]$.

3.5 Specifying the cost of sampling the dual vectors

It remains to specify the quantity $T_{\text{sampling}}(N, qs)$ appearing in Theorem 9. We proceed as in [35, Section 4], [37], [39, Section 4.3] and [36, Section 4.4] in order to make our estimates directly comparable with theirs, except that we use the exact sampler of Section 3.2 instead of the MCMC sampler of [40]. That is, the sampling process consists of the following steps: we first BKZ-reduce a basis of $L_q^{\perp}(\boldsymbol{A}_{\text{dual}})$ with block size β, which yields the basis $\boldsymbol{B} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m)$ of Section 3.1, determine the rank $t = t^*$ of Theorem 7 from its Gram–Schmidt norms, and then sample N vectors from $D_{\Lambda_t, qs}$ by the sampler of Section 3.2 on the basis $(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_t)$ of Λ_t. As this sampler calls Klein's algorithm $1/\Delta_t$ times per vector in expectation, we have

$$T_{\text{sampling}}(N, qs) \leq T_{\text{BKZ}}(m, \beta) + N \cdot \tfrac{1}{\Delta_t} \cdot \text{poly}(m), \tag{14}$$

where

$$\tfrac{1}{\Delta_t} \;=\; \frac{\prod_{i=1}^{t} \rho_{qs/\|\widetilde{\boldsymbol{b}}_i\|}(\mathbb{Z})}{\rho_{qs}(\Lambda_t)} \;\leq\; \prod_{i=1}^{t} \rho_{\|\widetilde{\boldsymbol{b}}_i\|/(qs)}(\mathbb{Z}). \tag{15}$$

For the inequality in (15) note that $\rho_{qs}(\Lambda_t) \geq (qs)^t / \det(\Lambda_t) = \prod_{i=1}^{t} qs/\|\widetilde{\boldsymbol{b}}_i\|$ by Lemma 1, and that $\rho_x(\mathbb{Z}) = x\, \rho_{1/x}(\mathbb{Z})$ for all $x > 0$. The bound (14) is the one of [36, Theorem 2] (cf. [39, Theorem 1]) for the MCMC sampler of [40], without the factor $\ln(1/\varepsilon)$, where ε is the statistical distance of the output of the latter from $D_{\Lambda_t, qs}$.

There remain two unknown quantities in (14): the Gram–Schmidt norms $\|\widetilde{\boldsymbol{b}}_i\|$ and $T_{\text{BKZ}}(m, \beta)$. We estimate both using the same formulas and cost model as in [39], so that our figures are directly comparable with theirs. That is, we predict the Gram–Schmidt norms by

$$\|\widetilde{\boldsymbol{b}}_i\| \;=\; H_\beta^{m-2i+1}\, q^{n_{\text{dual}}/m} \qquad \text{for } i = 1, \ldots, m, \tag{16}$$

where H_β denotes the root Hermite factor of BKZ with block size β, and we estimate $1/\Delta_t$ by the right-hand side of (15). Note that, both in ours and in

their setting, the lattice that is reduced has determinant $q^{n_{\mathrm{dual}}}$. The difference being that for us this holds with probability 1 (as the lower block of $\boldsymbol{A}_{\mathrm{dual}}$ is given by $\boldsymbol{I}_{n_{\mathrm{dual}}}$), and for them this holds with probability close to 1. What does differ, of course, is the dimension of this lattice: $\bar{m}$ for them, and $m = \bar{m} + n_{\mathrm{dual}}$ for us. Note that $T_{\mathrm{BKZ}}(m, \beta)$ is evaluated with the lattice estimator of [3] under the cost model of [2] with the sieving oracle of [10], which is the cost model used in [39, Section 4.3] and in [35,37].

4 A further improved attack

In Algorithm 1 every guessed coordinate is enumerated. In non-provable dual attacks it is common to treat a part of the guessed coordinates by a discrete Fourier transform instead; cf. [20]. A split of this kind is also used in the provable attack of [36], where the first part is guessed to be zero.

In this section we adapt this technique, but, here, the first part is enumerated in decreasing order of probability. We show that our analysis carries over without any new argument: the resulting algorithm differs from Algorithm 1 only in the order in which the candidates are examined.

The attack. Let $n_{\mathrm{guess}} = n_{\mathrm{enum}} + n_{\mathrm{fft}}$ with integers $n_{\mathrm{enum}}, n_{\mathrm{fft}} \geq 0$, and split

$$\bar{\boldsymbol{A}}_{\mathrm{guess}} = \left(\bar{\boldsymbol{A}}_{\mathrm{enum}}\ \bar{\boldsymbol{A}}_{\mathrm{fft}}\right), \qquad \boldsymbol{s}_{\mathrm{guess}} = \begin{pmatrix} \boldsymbol{s}_{\mathrm{enum}} \\ \boldsymbol{s}_{\mathrm{fft}} \end{pmatrix} \qquad \text{and} \qquad \tilde{\boldsymbol{s}}_{\mathrm{guess}} = \begin{pmatrix} \tilde{\boldsymbol{s}}_{\mathrm{enum}} \\ \tilde{\boldsymbol{s}}_{\mathrm{fft}} \end{pmatrix}$$

accordingly. Let R_{enum} and R_{fft} denote the rank functions of the distributions $\mathcal{C}(\eta)^{\otimes n_{\mathrm{enum}}}$ and $\mathcal{C}(\eta)^{\otimes n_{\mathrm{fft}}}$, defined as the rank function R in the proof of Theorem 7, and let $\textsc{GetRealization}_{\mathrm{enum}}$ and $\textsc{GetRealization}_{\mathrm{fft}}$ be the corresponding instances of $\textsc{GetRealization}$.

Algorithm 2 The improved dual attack

Input and output: As in Algorithm 1.
for $k \leftarrow 1$ **to** $(2\eta+1)^{n_{\mathrm{enum}}}$ **do**
  $\tilde{\boldsymbol{s}}_{\mathrm{enum}} \leftarrow \textsc{GetRealization}_{\mathrm{enum}}(k)$
  Compute $g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\mathrm{enum}}, \boldsymbol{x}))$ for all $\boldsymbol{x} \in \mathbb{Z}_q^{n_{\mathrm{fft}}}$ ▷ by one DFT; see (18).
  for $l \leftarrow 1$ **to** $(2\eta+1)^{n_{\mathrm{fft}}}$ **do**
    $\tilde{\boldsymbol{s}}_{\mathrm{fft}} \leftarrow \textsc{GetRealization}_{\mathrm{fft}}(l)$
    if $g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\mathrm{enum}}, \tilde{\boldsymbol{s}}_{\mathrm{fft}})) > c \cdot \mathcal{T}$ **then**
      return $(\tilde{\boldsymbol{s}}_{\mathrm{enum}}, \tilde{\boldsymbol{s}}_{\mathrm{fft}})$
    end if
  end for
end for
return $\perp$

Algorithm 2 tests the candidates $\tilde{\boldsymbol{s}}_{\mathrm{guess}} \in [\pm\eta]^{n_{\mathrm{guess}}}$ in increasing order of

$$R'(\tilde{\boldsymbol{s}}_{\mathrm{guess}}) \ := \ \big(R_{\mathrm{enum}}(\tilde{\boldsymbol{s}}_{\mathrm{enum}}) - 1\big)\,(2\eta+1)^{n_{\mathrm{fft}}} + R_{\mathrm{fft}}(\tilde{\boldsymbol{s}}_{\mathrm{fft}}), \tag{17}$$

which is a bijection from $[\pm\eta]^{n_{\text{guess}}}$ onto $\{1,\dots,(2\eta+1)^{n_{\text{guess}}}\}$, and returns the first one whose score exceeds $c\cdot\mathcal{T}$. That is, it coincides with Algorithm 1, except that the candidates are examined in the order given by R' instead of R. For $n_{\text{fft}}=0$ we return to Algorithm 1, and for $n_{\text{enum}}=0$ it is Algorithm 1 with all scores computed by one discrete Fourier transform; cf. the proof of Theorem 9.

The DFT-trick on $\mathbb{Z}^{n_{\text{fft}}}$. The scores in line 4 of Algorithm 2 are obtained as follows. We first extend (8) from $[\pm\eta]^{n_{\text{fft}}}$ to $\boldsymbol{x}\in\mathbb{Z}^{n_{\text{fft}}}$. Then, we note that $g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{enum}},\boldsymbol{x}))$ depends on $\boldsymbol{x}$ only modulo q. For $j=1,\dots,N$ let $\boldsymbol{w}_j^u\in\mathbb{Z}^{\bar m}$ consist of the first $\bar m$ coordinates of $\boldsymbol{w}_j$, and set $\boldsymbol{a}_j:=\bar{\boldsymbol{A}}_{\text{fft}}^{\mathsf{T}}\boldsymbol{w}_j^u \bmod q$ and $\varphi_j:=\langle \boldsymbol{b}-\bar{\boldsymbol{A}}_{\text{enum}}\tilde{\boldsymbol{s}}_{\text{enum}},\boldsymbol{w}_j^u\rangle$. Then, by (7) and (8),

$$\begin{aligned} g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\text{enum}},\boldsymbol{x})) &= \tfrac{1}{N}\operatorname{Re}\textstyle\sum_{\boldsymbol{a}\in\mathbb{Z}_q^{n_{\text{fft}}}}\Phi(\boldsymbol{a})\,\mathrm{e}^{-2\pi i\langle\boldsymbol{a},\boldsymbol{x}\rangle/q},\\ \text{where}\quad \Phi(\boldsymbol{a}) &:= \textstyle\sum_{j=1}^{N}\mathbb{1}_{(\boldsymbol{a}_j=\boldsymbol{a})}\,\mathrm{e}^{2\pi i\varphi_j/q}. \end{aligned} \tag{18}$$

Hence, line 4 amounts to computing Φ, which depends on $\tilde{\boldsymbol{s}}_{\text{enum}}$, and its discrete Fourier transform over $\mathbb{Z}_q^{n_{\text{fft}}}$. This costs $\operatorname{poly}(m,n)\cdot(N+q^{n_{\text{fft}}})$; cf. the proof of Theorem 9.

Theorem 10. *Let $(\bar{\boldsymbol{A}},\boldsymbol{b})$, m, Λ_t, c, γ_t and $\mathcal{T}$ be as in Theorem 7, let $n_{\text{guess}}=n_{\text{enum}}+n_{\text{fft}}$, and set $\mathcal{K}:=2^{n_{\text{enum}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}(2\eta+1)^{n_{\text{fft}}}$. Consider the inequality*

$$\tfrac{2}{c\mathcal{T}}\cdot m\,\mathcal{K}\,\rho_s(\mathbb{Z})^m\,\det(\Lambda_t)\;\le\;(qs)^t, \tag{19}$$

suppose that it holds for $t=m$ (which is a deterministic condition), and let t^* denote the smallest $t\in\{1,\dots,m\}$, for which (19) holds. Moreover, let

$$c\;=\;\tfrac{1}{2}\qquad\text{and}\qquad N\;:=\;\left\lceil\tfrac{8}{\mathcal{T}^2}\big(\ln(\mathcal{K})+\ln(m)\big)\right\rceil, \tag{20}$$

where $\mathcal{T}$ is evaluated at $t=t^*$ (so that N is a function of $\bar{\boldsymbol{A}}_{\text{dual}}$). Run Algorithm 2 with $t=t^*$ and with N dual vectors sampled as in Lemma 5. Then, Algorithm 2 is successful with probability at least $\frac{1}{2}-\frac{3}{m}$. If, moreover, $m>6$, then the expected complexity of Algorithm 2, conditioned on the event that it outputs $\boldsymbol{s}_{\text{guess}}$, is at most

$$\tfrac{1}{\frac{1}{2}-\frac{3}{m}}\mathbb{E}\Big[T_{\text{sampling}}(N,qs)+\operatorname{poly}(m,n)2^{n_{\text{enum}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}\big(N+q^{n_{\text{fft}}}\big)\Big]+\widetilde{\mathcal{O}}\big(n_{\text{guess}}^{2\eta}\big),$$

where the expectation is taken with respect to $\bar{\boldsymbol{A}}_{\text{dual}}$, as in Theorem 9.

Proof. The proof of Theorem 7 uses the rank function R in two ways only. First, it determines the order of the candidates, through the event (FP) and the events $(R(\tilde{\boldsymbol{s}}_{\text{guess}})<R(\boldsymbol{s}_{\text{guess}}))$, which are determined by $\boldsymbol{s}$. Second, as R is a bijection, $\sum_{\tilde{\boldsymbol{s}}_{\text{guess}}}\mathbb{E}[\mathbb{1}_{(R(\tilde{\boldsymbol{s}}_{\text{guess}})<R(\boldsymbol{s}_{\text{guess}}))}]=\mathbb{E}[R(\boldsymbol{s}_{\text{guess}})]-1$, and this is bounded by $2^{n_{\text{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}$ in the proofs of Lemma 17 and Lemma 20.

Hence, this proof applies to Algorithm 2, with (19), (20), R' and $\mathcal{K}$ in place of (11), (12), R and $2^{n_{\text{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}$, once we know that $\mathbb{E}[R'(\boldsymbol{s}_{\text{guess}})]\le\mathcal{K}$. The

latter holds due to the facts that $R_{\mathrm{fft}}(\boldsymbol{s}_{\mathrm{fft}}) \leq (2\eta+1)^{n_{\mathrm{fft}}}$ and $\mathbb{E}[R_{\mathrm{enum}}(\boldsymbol{s}_{\mathrm{enum}})] \leq 2^{n_{\mathrm{enum}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}$ by Arikan's inequality (cf. the proof of Theorem 9).

For the complexity we argue as in the proof of Theorem 9, where the term $\widetilde{\mathcal{O}}(n_{\mathrm{guess}}^{2\eta})$ now accounts for the pre-computation phases of both instances of GetRealization. Let S be the event that Algorithm 2 outputs $\boldsymbol{s}_{\mathrm{guess}}$. On S the outer loop performs $R_{\mathrm{enum}}(\boldsymbol{s}_{\mathrm{enum}})$ iterations. Each iteration costs $\mathrm{poly}(m,n)\cdot(N+q^{n_{\mathrm{fft}}})$: this is the cost of line 4 (computed above), and the inner loop performs at most $(2\eta+1)^{n_{\mathrm{fft}}} \leq q^{n_{\mathrm{fft}}}$ iterations. As $\boldsymbol{s}_{\mathrm{enum}}$ is independent of $\bar{\boldsymbol{A}}_{\mathrm{dual}}$, the conditional expectation of $R_{\mathrm{enum}}(\boldsymbol{s}_{\mathrm{enum}})$ given $\bar{\boldsymbol{A}}_{\mathrm{dual}}$ is at most $2^{n_{\mathrm{enum}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}$. As in Theorem 9, conditioning on S contributes the factor $(\frac{1}{2}-\frac{3}{m})^{-1}$.

5 Application to ML-KEM

The goal of this section is to run Algorithm 2 for the three parameter sets of ML-KEM ([33]) and to apply Theorem 10. Hereby, we determine, for each of the parameter sets, choices of the free parameters, which minimize the resulting complexity bound in some sense.

5.1 On the module structure of ML-KEM

ML-KEM is a *module-LWE (mLWE)* scheme. However, we do not exploit the module structure here in this work, so that we treat ML-KEM as an LWE instance by viewing the module of rank k over a ring of degree 256 as a lattice problem in dimension $n = 256k$.

Violation of the uniformness of $\bar{\boldsymbol{A}}$. This simplification has some theoretic consequences. More precisely, recall that in Definition 3 the matrix $\bar{\boldsymbol{A}}$ is (in particular) assumed to be uniformly distributed on $\mathbb{Z}_q^{\bar{m}\times n}$. In ML-KEM the matrix is generated differently:

- The matrix is not sampled uniformly at random, but is expanded deterministically from a public seed by an extendable output function.
- Moreover, viewed as a matrix over $\mathbb{Z}_q$, for ML-KEM, the matrix $\bar{\boldsymbol{A}}$ is a $k\times k$ block matrix, whose blocks are determined by 256 ring coefficients each. It therefore has only $256k^2$ degrees of freedom instead of $\bar{m}n = (256k)^2$. In particular, $\bar{\boldsymbol{A}}$ is *not* uniformly distributed on $\mathbb{Z}_q^{\bar{m}\times n}$. Moreover, as each block is determined by any one of its columns, and as $n_{\mathrm{guess}} < 256$ for all parameters below, $\bar{\boldsymbol{A}}_{\mathrm{guess}}$ is even a deterministic function of $\bar{\boldsymbol{A}}_{\mathrm{dual}}$. Hence the conclusion of Lemma 18 fails for every realization of $\bar{\boldsymbol{A}}_{\mathrm{dual}}$.

The first deviation is the usual assumption that the extendable output function is treated as a random oracle. This is a typical assumption made throughout the literature on the concrete security of ML-KEM. The second deviation is much more critical and it is the main reason why the hypothesis of Definition 3 does not hold for ML-KEM.

Formulation of the heuristic assumption. Fortunately, in our proofs, the hypothesis on $\bar{\boldsymbol{A}}$ of Definition 3 is used *at exactly one point*, namely for Lemma 19 (via Lemma 18), and this lemma is only used in order to derive the first inequality of (34) in Lemma 20. Every other step of Section 8 and Section 9 uses only Lemma 5 (whose proof does not use the distribution of $\bar{\boldsymbol{A}}$), the independence of $\bar{\boldsymbol{A}}$ and $(\boldsymbol{s}, \boldsymbol{e})$, and deterministic properties of $\boldsymbol{A}_{\mathrm{dual}}$ and of the reduced basis. We therefore assume this inequality directly.

Heuristic 1 *Let $\bar{\boldsymbol{A}}$ be generated as in ML-KEM [33] and let $\boldsymbol{s}$ and $\boldsymbol{e}$ be as in Definition 3 and independent of $\bar{\boldsymbol{A}}$. Let c, $t = t^*$, $\mathcal{T}$, W and $\mathcal{K}$ be as in Theorem 10, and let R' be as in* (17)*. Then*

$$\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}} \in [\pm\eta]^{n_{\mathrm{guess}}}} \mathbb{E}\Big[\mathbb{1}_{(R'(\tilde{\boldsymbol{s}}_{\mathrm{guess}}) < R'(\boldsymbol{s}_{\mathrm{guess}}))} \cdot \tfrac{2}{c\mathcal{T}} \cdot \mathbb{E}[g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\mathrm{guess}})) \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] \Big] \leq \tfrac{1}{m\mathcal{K}} \left(\mathbb{E}\left[R'(\boldsymbol{s}_{\mathrm{guess}}) \right] - 1 \right).$$

Consequently, by additionally assuming Heuristic 1, the claims of Theorem 10 remain unchanged. Intuitively, Heuristic 1 asserts that, on average over $\bar{\boldsymbol{A}}_{\mathrm{dual}}$, $\boldsymbol{s}$ and $\boldsymbol{e}$, the expected scores of the wrong candidates do not exceed the bound which Lemma 19 and (19) yield for a uniformly distributed matrix. It is left for future research to rigorously verify this heuristic for ML-KEM.

5.2 The non-free parameters

As explained above, in ML-KEM, $\bar{\boldsymbol{A}}$ is a matrix in $\mathbb{Z}_q^{\bar{m} \times n}$ with $\bar{m} = n = 256 \cdot k$, where $k = 2, 3, 4$ for ML-KEM-512, ML-KEM-768 and ML-KEM-1024, respectively. Moreover, the key generation in ML-KEM produces $\boldsymbol{b} = \bar{\boldsymbol{A}}\boldsymbol{s} + \boldsymbol{e}$ with both $\boldsymbol{s}$ and $\boldsymbol{e}$ drawn from $\mathcal{C}(\eta_1)$, so that Definition 3 applies with $\eta = \eta_1$. In all three cases $q = 3329$, while $\eta = 3$ for ML-KEM-512 and $\eta = 2$ for the other two sets. We have $\mathcal{H}_{1/2}(\mathcal{C}(3)) = 2.5136$ and $\mathcal{H}_{1/2}(\mathcal{C}(2)) = 2.1577$.

5.3 The free parameters

After fixing $q, n, \bar{m}$ and η (as in Section 5.2), the bounds of Theorem 10 depend on four free parameters, namely n_{enum}, n_{fft}, s and the block size β of the BKZ reduction in Section 3.5. The remaining quantities are then determined: n_{guess} (by $n_{\mathrm{enum}} + n_{\mathrm{fft}}$), n_{dual} (by $n - n_{\mathrm{guess}}$), m (by (4)), c (by (20)), t^* (by the reduced basis, cf. Theorem 10), and, for $t = t^*$, γ_t and $\mathcal{T}$ (by (10)) and N (by (20)).

Admissible parameters. The hypothesis of Theorem 10 is that (19) holds for $t = m$. As $\det(\Lambda_m) = q^{n_{\mathrm{dual}}}$, this is an explicit condition on $(n_{\mathrm{enum}}, n_{\mathrm{fft}}, s)$; cf. Remark 8. We call a quadruple $(n_{\mathrm{enum}}, n_{\mathrm{fft}}, s, \beta)$ *admissible* if this condition is satisfied.

The cost. For $t \in \{1, \dots, m\}$ let

$$T(t) = T_{\mathrm{BKZ}}(m, \beta) + N \prod_{i=1}^{t} \rho_{\|\widetilde{\boldsymbol{b}}_i\|/(qs)}(\mathbb{Z}) + 2^{n_{\mathrm{enum}} \mathcal{H}_{1/2}(\mathcal{C}(\eta))} \big(N + q^{n_{\mathrm{fft}}}\big), \quad (21)$$

where N is given by (20) for this t, where the product is the right-hand side of (15), and where we suppress the factors $\mathrm{poly}(m, n)$ and $(\frac{1}{2} - \frac{3}{m})^{-1}$ and the additive term $\widetilde{\mathcal{O}}(n_{\mathrm{guess}}^{2\eta})$. By Theorem 10, (14) and (15), the expected complexity of our attack is then at most $\mathbb{E}[T(t^*)]$. As N and the right-hand side of (15) are non-decreasing in t, so is T, and every upper bound on t^* yields an upper bound on $\mathbb{E}[T(t^*)]$. The following two subsections differ in the upper bound on t^* that is used.

The search. In both cases we minimize T by an exhaustive search over all admissible quadruples $(n_{\mathrm{enum}}, n_{\mathrm{fft}}, s, \beta)$ with $n_{\mathrm{enum}}, n_{\mathrm{fft}} \geq 0$, $n_{\mathrm{enum}} + n_{\mathrm{fft}} \leq n$ and $60 \leq \beta \leq m$. For every $(n_{\mathrm{enum}}, n_{\mathrm{fft}})$, the parameter s runs over a geometric grid of 400 values between a lower and an upper bound on s (see below), which we refine tenfold around the optimum (n_{guess} within 8 for all n_{fft}, β within 16). The results are thus optimal in n_{enum}, n_{fft}, β and t, and in s up to the grid resolution. BKZ requires $\beta \leq m$. Block sizes below 60 are not considered, as a smaller block size only deteriorates the profile (16), and as already at $\beta = 60$ the smallest cost found on a coarse grid (about $2^{823.8}$, $2^{1356.7}$ and $2^{2053.5}$) exceeds the costs in Table 1 and Table 2 by more than 600 bits.

The range of values for s is deduced from two inequalities. First, admissibility means that (19) holds with $\det(\Lambda_t)$ replaced by $q^{n_{\mathrm{dual}}}$ (in Section 5.4) or by the right-hand side of (22) (in Section 5.5). Both quantities are at least $q^{t n_{\mathrm{dual}}/m}$. As $\frac{2}{c\mathcal{T}} \geq 1$ and $\rho_s(\mathbb{Z}) \geq 1$, this implies $m\mathcal{K} \leq \big(s\, q^{\bar{n}/m}\big)^t$, which yields a lower bound on s. Second, as the last summand of (21) is at least N, so is $T(t)$, and, by (20) and (10), $\log(N) \geq \log\big(8(\ln(\mathcal{K}) + \ln(m))\big) + \frac{\pi s^2 \eta\,(t + 2\sqrt{t\ln(2)} + 2\ln(2))}{\ln(2)}$. Hence, large values of s lead to a large cost, which yields an upper bound on s. For all three parameter sets and all searches, the resulting grids lie in $[0.001, 3.2]$.

A Python script that reproduces every entry of our rows of Table 1, Table 2 and Table 3 and reruns all searches (including the comparison in Remark 11) is available at `https://codeberg.org/BashiriK/Further_Elaborations_Dual_Attack`.

5.4 Minimization for the full lattice

Here we only use $t^* \leq m$, that is, we bound $\mathbb{E}[T(t^*)]$ by $T(m)$, the cost for the full lattice $L_q^{\perp}(\boldsymbol{A}_{\mathrm{dual}})$, and we minimize $T(m)$ over all admissible quadruples. The reduced basis then enters only through the Gram–Schmidt norms in (21), which we predict by (16) exactly as in [35,37,39], where the dual vectors are likewise sampled from a full lattice. The result is given in Table 1: our attack improves upon [35] by 21.6, 36.2 and 52.2 bits. The concurrent work [29] reports 2^{216}, 2^{317} and 2^{426}.

ML-KEM		n_{guess}	n_{fft}	n_{dual}	m	β	s	$\log(N)$	T_{BKZ}	total
512	[35]	19	19	493	1023	746	0.090	n.c.	n.c.	2^{238}
	this work	58	7	454	966	667	0.0740	85.7	$2^{215.5}$	$2^{216.4}$
768	[35]	29	29	739	1519	1129	0.110	n.c.	n.c.	2^{347}
	this work	98	10	670	1438	1003	0.0890	118.2	$2^{310.1}$	$2^{310.8}$
1024	[35]	39	39	985	2025	1602	0.110	n.c.	n.c.	2^{478}
	this work	138	13	886	1910	1411	0.0891	153.5	$2^{424.8}$	$2^{425.8}$

Table 1. Optimised parameters for Algorithm 2 for the full lattice (that is, $t = m$), compared with [35] as reported in [39, Table 1]. The attacks of [37,39] are more expensive than [35] and therefore not included. Entries n.c. are not communicated in [39, Table 1]. In [35], we have $n_{\text{fft}} = n_{\text{guess}}$, and m is the number of LWE samples used there (the dimension of their lattice), which exceeds the n samples provided by ML-KEM. In our rows, $n_{\text{enum}} = n_{\text{guess}} - n_{\text{fft}}$, and total is the cost $T(m)$ of (21). By Theorem 10, our success probability is at least $\frac{1}{2} - \frac{3}{m}$, and no assumption on the shape of the reduced basis enters this bound.

5.5 Optimization under the geometric series assumption

In order to be as comparable as possible with [36], we treat $\det(\Lambda_t)$ as they do (see [36, Lemma 6]), but we allow for a factor $\kappa \geq 1$. For our purposes the following assumption suffices.

Heuristic 2 *Let $\boldsymbol{B}$ be the reduced basis of Section 3.5, let $\kappa \geq 1$, and let $t \in \{1, \ldots, m\}$. Then, for every realization of $\bar{\boldsymbol{A}}_{\text{dual}}$,*

$$\det(\Lambda_t) \;\leq\; \kappa \prod_{i=1}^{t} H_\beta^{m-2i+1}\, q^{n_{\text{dual}}/m} \;=\; \kappa\, H_\beta^{t(m-t)}\, q^{t n_{\text{dual}}/m}. \tag{22}$$

For $\kappa = 1$, the right-hand side of (22) is the value predicted by the profile (16), as in [36, Lemma 6]. Larger values of κ weaken Heuristic 2. As Table 2 shows, our costs depend only mildly on κ: compared with $\kappa = 1$, they increase by at most 1.2 bits for $\kappa = 2^{32}$, and by at most 3.9 bits for $\kappa = 2^{128}$. We call a tuple $(n_{\text{enum}}, n_{\text{fft}}, s, \beta, t)$ *admissible* if it satisfies (19) with $\det(\Lambda_t)$ replaced by the right-hand side of (22). As the left-hand side of (19) is increasing in $\det(\Lambda_t)$, Heuristic 2 implies that (19) then holds for t itself (if t belongs to an admissible tuple). By the definition of t^*, this gives $t^* \leq t$, and hence $\mathbb{E}[T(t^*)] \leq \mathbb{E}[T(t)]$. We therefore minimize $T(t)$ over all admissible tuples. Taking logarithms in (19) and inserting (10) and the right-hand side of (22) shows that $(n_{\text{enum}}, n_{\text{fft}}, s, \beta, t)$ is admissible if and only if $F(t) \leq 0$, where

$$\begin{aligned} F(t) \;:=\; & \ln\left(\tfrac{2}{c}\,\kappa\, m\, \mathcal{K}\, \rho_s(\mathbb{Z})^m\right) + \pi s^2 \tfrac{\eta}{2}\left(t + 2\sqrt{t \ln(2)} + 2\ln(2)\right) \\ & +\; t(m-t)\ln(H_\beta) - t \ln\left(s\, q^{\bar{m}/m}\right). \end{aligned}$$

As $H_\beta \geq 1$, the function F is concave on $(0, \infty)$, and $F(t) > 0$ for $t \to 0$. Hence the admissible values of t form a set of consecutive integers ending at m if $(n_{\text{enum}}, n_{\text{fft}}, s, \beta, m)$ is admissible, and otherwise there are none. In the former

case, the quadruple $(n_{\text{enum}}, n_{\text{fft}}, s, \beta)$ is admissible as well, as $\det(\Lambda_m) = q^{n_{\text{dual}}} \leq \kappa\, q^{n_{\text{dual}}}$ (for $\kappa = 1$, the two conditions are equivalent). As T is non-decreasing in t, the optimal t is the smallest admissible one, which we determine by bisection. Under Heuristic 2 it is an upper bound on t^*. The result is given in Table 2: for $\kappa = 1$, compared with Table 1, the bound on the rank of the sublattice saves about 25.7, 35.9 and 47.6 bits.

ML-KEM		κ	n_{guess}	n_{fft}	n_{dual}	m	t	β	s	$\log(N)$	T_{BKZ}	total
512	[36]	1	17	16	495	1185	560	650	0.134	152.9	n.c.	$2^{211.9}$
	this work	1	44	8	468	980	379	574	0.1252	97.4	$2^{190.0}$	$2^{190.7}$
	this work	2^{32}	45	8	467	979	398	578	0.1214	96.1	$2^{191.1}$	$2^{191.8}$
	this work	2^{128}	46	8	466	978	445	588	0.1148	95.8	$2^{193.8}$	$2^{194.5}$
768	[36]	1	27	25	741	1551	810	1020	0.171	232.0	n.c.	$2^{317.1}$
	this work	1	76	11	692	1460	541	874	0.1521	131.8	$2^{274.3}$	$2^{274.9}$
	this work	2^{32}	77	11	691	1459	559	877	0.1492	131.0	$2^{275.2}$	$2^{275.9}$
	this work	2^{128}	78	11	690	1458	605	887	0.1436	131.0	$2^{277.9}$	$2^{278.6}$
1024	[36]	1	37	35	987	2167	1090	1440	0.169	300.5	n.c.	$2^{434.3}$
	this work	1	111	14	913	1937	746	1241	0.1470	165.8	$2^{377.5}$	$2^{378.1}$
	this work	2^{32}	111	14	913	1937	760	1244	0.1465	167.5	$2^{378.4}$	$2^{379.1}$
	this work	2^{128}	112	14	912	1936	805	1254	0.1425	167.6	$2^{381.1}$	$2^{381.8}$

Table 2. Under Heuristic 2 (that is, $t^* \leq t$): optimised parameters for Algorithm 2, compared with [36, Table 1], which uses Heuristic 2 with $\kappa = 1$ in the same way. Recall that the factor κ enters only through the estimation of t^*. In all rows, $1/\Delta_t$ is estimated by (16), as in Table 1. In the rows of [36], n_{guess} is the number of coordinates of the secret that are not part of the lattice, of which $n_{\text{guess}} - n_{\text{fft}}$ are guessed to be zero, and m is the dimension of the lattice; cf. Remark 12. Entries n.c. are not communicated in [36, Table 1]. In our rows, $n_{\text{enum}} = n_{\text{guess}} - n_{\text{fft}}$, and total is the cost $T(t)$ of (21). Our success probability is at least $\frac{1}{2} - \frac{3}{m}$, which is derived without Heuristic 2.

Remark 11. (Advantage due to the guessing strategy) The values of n_{guess} in both tables are much larger than the numbers of coordinates of the secret that are not part of the lattice in [35] and [36]. Here we benefit from the guessing strategy of Section 3.1: by Theorem 10, each enumerated coordinate costs, in expectation, a factor of at most $2^{\mathcal{H}_{1/2}(\mathcal{C}(\eta))}$, that is, $2^{2.5136}$ respectively $2^{2.1577}$, whereas the discrete Fourier transform over $\mathbb{Z}_q$ costs $q \approx 2^{11.7}$ per coordinate once $q^{n_{\text{fft}}}$ exceeds N. Each guessed coordinate decreases n_{dual}, and thereby the dimension m of the lattice that has to be reduced. Under Heuristic 2 with $\kappa = 1$, the best parameters we found with $n_{\text{enum}} = 0$ (i.e. if all guessed coordinates are treated by the Fourier transform) have $n_{\text{guess}} = 16$, 24 and 34, and cost $2^{198.5}$, $2^{290.4}$ and $2^{402.0}$ (for $n_{\text{enum}} = 0$, Algorithm 2 is Algorithm 1, so that Theorem 7 applies). The enumeration thus accounts for an improvement of about 7.8, 15.5 and 23.8 bits.

Remark 12. (Comparison with [36]) Our analysis differs from [36] as follows. First, their parameters use 690, 810 and 1180 LWE samples, whereas we use the $\bar{m} = n$ samples that ML-KEM provides. Second, in [36] the wrong candidates are controlled through the lengths of the shortest nonzero vectors of two projected lattices. These lengths are estimated by the Gaussian heuristic and cannot be verified by the attacker. In our analysis of the wrong candidates only $\det(\Lambda_t)$ enters, and only through the cost (via t^*). Moreover, (22) can be verified by the attacker once the basis is reduced. Third, in [36] the projected error is replaced by its typical length, whereas the term γ_t in (10) accounts for its deviation, so that our bound on the success probability is explicit. Fourth, in [36] a further part of the secret is guessed to be zero, in 3, 7 and 7 independent repetitions, and this guess is correct in at least one of them with probability 0.649, 0.608 and 0.620; cf. [36, Lemma 11 and Table 1]. Their attack thus succeeds with at most this probability, whereas ours succeeds with probability at least $\frac{1}{2} - \frac{3}{m}$.

6 Application to ML-DSA

In ML-DSA [32], the key generation computes $\boldsymbol{t} = \bar{\boldsymbol{A}}\boldsymbol{s}_1 + \boldsymbol{s}_2$ over $\mathcal{R}_q$, where $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^{256}+1)$ with $q = 8380417$, $\bar{\boldsymbol{A}} \in \mathcal{R}_q^{k\times\ell}$ is expanded from a public seed, and where the coefficients of $\boldsymbol{s}_1 \in \mathcal{R}_q^{\ell}$ and $\boldsymbol{s}_2 \in \mathcal{R}_q^{k}$ are drawn independently from $\mathcal{U}(\eta)$, the uniform distribution on $\{-\eta, \ldots, \eta\}$. The public key contains only the high-order part $\boldsymbol{t}_1$ of $\boldsymbol{t} = 2^{13}\boldsymbol{t}_1 + \boldsymbol{t}_0$. As usual in the security analysis of ML-DSA, which relies on the hardness of MLWE for the pair $(\bar{\boldsymbol{A}}, \boldsymbol{t})$ [7, Section 6.2], we assume that the full $\boldsymbol{t}$ is known. Viewed over $\mathbb{Z}_q$, the pair $(\bar{\boldsymbol{A}}, \boldsymbol{t})$ is then an LWE instance with $n = 256\ell$ and $\bar{m} = 256k$, secret $\boldsymbol{s} = \boldsymbol{s}_1$ and error $\boldsymbol{e} = \boldsymbol{s}_2$, where $(k, \ell, \eta) = (4, 4, 2)$, $(6, 5, 4)$ and $(8, 7, 2)$ for ML-DSA-44, ML-DSA-65 and ML-DSA-87, respectively.

Adaptations. This instance differs from Definition 3 in the distribution $\mathcal{U}(\eta)$ instead of $\mathcal{C}(\eta)$, and, as for ML-KEM, in the structure of $\bar{\boldsymbol{A}}$. With minimal adaptations of Definition 3 and of the proofs, our results carry over. Indeed, the distribution of $\boldsymbol{s}$ and $\boldsymbol{e}$ enters our analysis only at the following three points.

First, Lemma 16 uses the sub-Gaussian parameter $\frac{\eta}{2}$ of $\mathcal{C}(\eta)$. The distribution $\mathcal{U}(\eta)$ is sub-Gaussian with parameter $\sigma^2 := \frac{\eta(\eta+1)}{3}$, its variance. Indeed, by [4, Proposition 3.3], as $\mathcal{U}(\eta)$ is symmetric, it suffices that $\mathbb{E}[\xi^{2k}] \le (2k-1)!!\,\sigma^{2k}$ for all $k \ge 2$ and $\xi \sim \mathcal{U}(\eta)$. This holds by direct computation for $k \in \{2, 3\}$ and $\eta \in \{2, 4\}$, and for $k \ge 4$ as $\mathbb{E}[\xi^{2k}] \le \eta^{2k} \le (3\sigma^2)^k \le (2k-1)!!\,\sigma^{2k}$, where we used $|\xi| \le \eta$, $\sigma^2 \ge \frac{\eta^2}{3}$ and $(2k-1)!! \ge 3^k$. Hence, Lemma 16, and all results based on it, hold with σ^2 in place of $\frac{\eta}{2}$, in particular in (10).

Second, Arikan's inequality yields the bound $2^{n_{\mathrm{enum}}\mathcal{H}_{1/2}(\mathcal{U}(\eta))} = (2\eta+1)^{n_{\mathrm{enum}}}$ in place of $2^{n_{\mathrm{enum}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}$, so that $\mathcal{K} = (2\eta+1)^{n_{\mathrm{guess}}}$.

Third, Lemma 18 uses only that $2\eta < q$ and that q is prime, both of which hold. Finally, as for ML-KEM, $\bar{\boldsymbol{A}}$ is not uniformly distributed on $\mathbb{Z}_q^{\bar{m}\times n}$, and we assume the analogue of Heuristic 1, with $\bar{\boldsymbol{A}}$ generated as in ML-DSA and with

$\mathcal{U}(\eta)$ in place of $\mathcal{C}(\eta)$. The verification of this heuristic for ML-DSA is left for future research.

Results. We proceed exactly as in Section 5, with σ^2 in place of $\frac{\eta}{2}$, with $\log(2\eta+1)$ in place of $\mathcal{H}_{1/2}(\mathcal{C}(\eta))$, and with Heuristic 2 for $\kappa = 1$. Table 3 reports the optimised parameters.

	n_{guess}	n_{fft}	n_{dual}	m	t	β	s	$\log(N)$	T_{BKZ}	total
ML-DSA-44	36	4	988	2012	329	508	0.1149	95.3	$2^{173.5}$	$2^{173.7}$
ML-DSA-65	38	6	1242	2778	451	766	0.0667	140.6	$2^{245.9}$	$2^{246.2}$
ML-DSA-87	81	6	1711	3759	635	1050	0.1060	148.3	$2^{325.9}$	$2^{326.1}$

Table 3. Optimised parameters for Algorithm 2 under Heuristic 2 with $\kappa = 1$ (assuming the full $\boldsymbol{t}$ to be known). The success probability is at least $\frac{1}{2} - \frac{3}{m}$.

7 Experimental demonstration

The complexity of our attacks against the ML-KEM parameters is, of course, too high for an implementation. In this section we therefore verify Theorem 7 practically, by executing Algorithm 1 on small instances and comparing the empirical frequency of success with the guaranteed lower bound $\frac{1}{2} - \frac{3}{m}$ from Theorem 7. The corresponding code can be found in `https://codeberg.org/BashiriK/Further_Elaborations_Dual_Attack`.

For each of the parameter sets of Table 4 we first verify the hypothesis of Theorem 7, that is, (11) for $t = m$. We then sample $\bar{\boldsymbol{A}}$, $\boldsymbol{s}$ and $\boldsymbol{e}$ according to Definition 3, reduce a basis of $L_q^{\perp}(\boldsymbol{A}_{\text{dual}})$ by LLL [27], and determine t^* from the Gram–Schmidt norms of the reduced basis. That is, we determine the smallest t for which (11) holds with the true value of $\det(\Lambda_t)$. Finally, we run Algorithm 1 with $t = t^*$, and with c and N as in (12), and record whether it extracts $\boldsymbol{s}_{\text{guess}}$. According to Theorem 7, the attack succeeds with probability at least $\frac{1}{2} - \frac{3}{m}$. In the first five parameter sets we chose $s = 0.12$. This is larger than necessary, as (11) holds for $t = m$ for all s between 0.088 and 1.0. We chose it such that $t^* < m$, and such that t^* varies between the instances. The last parameter set illustrates a larger dimension ($m = 136$). Here, $q = 2053$ keeps the number $1/\Delta_t$ of calls to Klein's algorithm per dual vector small, and we chose $s = 0.17$ ((11) holds for $t = m$ for all s between 0.018 and 1.4). One instance takes about 17 minutes on average on two cores.

The dual vectors are sampled exactly from $D_{\Lambda_t, qs}$, by the sampler of Section 3.2 on the first $t = t^*$ vectors of the LLL-reduced basis of $L_q^{\perp}(\boldsymbol{A}_{\text{dual}})$.

n	n_{guess}	m	s	t^*	N	$\frac{1}{2} - \frac{3}{m}$	success
11	6	16	0.12	11–12	166–175	0.312	1000/1000
15	7	23	0.12	16–18	253–282	0.369	500/500
19	8	30	0.12	22–23	393–414	0.400	300/300
25	10	40	0.12	30–32	720–798	0.425	100/100
31	12	50	0.12	38–40	1267–1403	0.440	40/40
75	14	136	0.17	45–52	29589–60279	0.477	6/6

Table 4. Experimental verification. In all cases $\eta = 1$, $\bar{m} = n$ and $n_{\text{dual}} = n - n_{\text{guess}}$; moreover, $q = 97$ in the first five rows and $q = 2053$ in the last row. The values of $\frac{1}{2} - \frac{3}{m}$ are rounded down. The columns t^* and N report the smallest and the largest value that occurred among the instances. The last column reports the number of instances in which Algorithm 1 returned $\boldsymbol{s}_{\text{guess}}$, out of the number of independent instances.

Remark 13. The bound of Theorem 7 holds in all six cases, and is conservative: Algorithm 1 returned $\boldsymbol{s}_{\text{guess}}$ in every instance. Note also that the rank t^* indeed depends on the instance.

In order to confirm that the experiment is not trivially successful, we repeated the first parameter set with N reduced below the value prescribed by (12). For $N = 60$, $N = 15$ and $N = 4$ the attack returned $\boldsymbol{s}_{\text{guess}}$ in 98, 18 and 6 out of 100 instances, respectively. The value of N in (12) is thus conservative, but the agreement in Table 4 is due to a sufficiently large N and not to the attack succeeding regardless of it.

8 The false-negative event

In this section we upper bound the probability of the false-negative event (FN), which we introduced in the proof of Theorem 7. First, in Section 8.1, we prove a bound on the conditional expectation of $g_W(\boldsymbol{t}(\boldsymbol{s}_{\text{guess}}))$ given $(\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})$. Then, using this bound, we prove the upper bound in Section 8.2. We conclude this section with a concentration result of the projected error vector.

8.1 Estimating the conditional expectation

Lemma 14. *Suppose the same setting as in Theorem 7, with $t := t^*$, except that (12) is not assumed, that is, $c \in (0,1)$ and $N \in \mathbb{N}$ are arbitrary, where N may be a function of $\bar{\boldsymbol{A}}_{\text{dual}}$ (and where (11), and thereby t^*, is understood with this c). Moreover, let $\boldsymbol{t}(\boldsymbol{s}_{\text{guess}})$ be defined as in (8) by choosing $\tilde{\boldsymbol{s}}_{\text{guess}} = \boldsymbol{s}_{\text{guess}}$. Then,*

$$\mathbb{E}\left[g_W(\boldsymbol{t}(\boldsymbol{s}_{\text{guess}})) \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})\right] \;\geq\; \rho_{\frac{1}{s}}\big(\operatorname{dist}\big(\pi\big(\begin{pmatrix}\boldsymbol{e}\\ -\boldsymbol{s}_{\text{dual}}\end{pmatrix}\big), \pi(L_q(\boldsymbol{A}_{\text{dual}}))\big)\big) \tag{23}$$

where π denotes the orthogonal projection onto $V = \operatorname{span}_{\mathbb{R}}(\Lambda_t)$; cf. (5).

Proof. The proof uses the same strategy as in the previous works on this topic; see e.g. [35, Section 4.2] or [9, Section 4.2]. In order to simplify the notation in this proof, we write $\boldsymbol{t} := \boldsymbol{t}(\boldsymbol{s}_{\text{guess}})$.

Note that $\mathbb{E}\left[g_W(\boldsymbol{t}) \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})\right] = \frac{1}{N}\sum_{i=1}^{N} \mathbb{E}[\cos(2\pi\frac{1}{q}\langle \boldsymbol{t}, \boldsymbol{w_i}\rangle) \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})]$. We will show that each summand is uniformly bounded from below by the right-hand side in (23). We therefore discard the dependency on i in the notation and abbreviate $\boldsymbol{w} := \boldsymbol{w}_i$. We also write $\boldsymbol{w} = \left((\boldsymbol{w}^u)^\mathsf{T}, (\boldsymbol{w}^d)^\mathsf{T}\right)^\mathsf{T}$ for $\boldsymbol{w}^u \in \mathbb{Z}^{\bar{m}}$ and $\boldsymbol{w}^d \in \mathbb{Z}^{n_{\mathrm{dual}}}$. Then, as $\boldsymbol{b} = \bar{\boldsymbol{A}}\boldsymbol{s} + \boldsymbol{e} = \bar{\boldsymbol{A}}_{\mathrm{dual}}\boldsymbol{s}_{\mathrm{dual}} + \bar{\boldsymbol{A}}_{\mathrm{guess}}\boldsymbol{s}_{\mathrm{guess}} + \boldsymbol{e} \bmod q$,

$$\begin{aligned}
\mathbb{E}[\cos(2\pi\tfrac{1}{q}\langle \boldsymbol{t}, \boldsymbol{w}\rangle) \,|\, (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] &= \mathbb{E}[\cos\left(2\pi\tfrac{1}{q}\left\langle \begin{pmatrix} \boldsymbol{b} - \bar{\boldsymbol{A}}_{\mathrm{guess}}\boldsymbol{s}_{\mathrm{guess}} \\ 0 \end{pmatrix}, \boldsymbol{w}\right\rangle\right) |(\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] \\
&= \mathbb{E}[\cos(2\pi\tfrac{1}{q}(\langle \bar{\boldsymbol{A}}_{\mathrm{dual}}\boldsymbol{s}_{\mathrm{dual}}, \boldsymbol{w}^u\rangle + \langle \boldsymbol{e}, \boldsymbol{w}^u\rangle)) \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] \\
&= \mathbb{E}[\cos(2\pi\tfrac{1}{q}(-\langle \boldsymbol{s}_{\mathrm{dual}}, \boldsymbol{w}^d\rangle + \langle \boldsymbol{e}, \boldsymbol{w}^u\rangle)) \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] \\
&= \mathbb{E}\left[\cos\left(2\pi\tfrac{1}{q}\left\langle \begin{pmatrix} \boldsymbol{e} \\ -\boldsymbol{s}_{\mathrm{dual}} \end{pmatrix}, \boldsymbol{w}\right\rangle\right) \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})\right],
\end{aligned} \tag{24}$$

where we used in the second-last step that $\frac{\boldsymbol{w}}{q} \in \frac{1}{q}L_q^\perp(\boldsymbol{A}_{\mathrm{dual}})$, that $\frac{1}{q}L_q^\perp(\boldsymbol{A}_{\mathrm{dual}})$ is the dual lattice to $L_q(\boldsymbol{A}_{\mathrm{dual}})$, and that $x \mapsto \cos(2\pi x/q)$ is invariant under adding integer multiples of q. We abbreviate $\boldsymbol{\xi} := \left(\boldsymbol{e}, -\boldsymbol{s}_{\mathrm{dual}}\right)^\mathsf{T} \in \mathbb{Z}^m$.

Note that $\cos(2\pi\frac{1}{q}\langle \boldsymbol{\xi}, \boldsymbol{w}\rangle)$ is of the form $G(\boldsymbol{w}, \bar{\boldsymbol{A}}, (\boldsymbol{s}, \boldsymbol{e}))$ for a bounded measurable function G, and that, by Lemma 5, the variables $\boldsymbol{w}$ and $(\boldsymbol{s}, \boldsymbol{e})$ are independent with respect to $\mathbb{P}[\cdot \mid \bar{\boldsymbol{A}}]$. Hence Lemma 23, applied with $\boldsymbol{\Gamma} = \bar{\boldsymbol{A}}$, $\boldsymbol{\Theta} = (\boldsymbol{s}, \boldsymbol{e})$ and $\boldsymbol{X} = \boldsymbol{w}$, yields

$$\mathbb{E}[\cos(2\pi\tfrac{1}{q}\langle \boldsymbol{\xi}, \boldsymbol{w}\rangle) \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] = \mathbb{E}[\cos(2\pi\tfrac{1}{q}\langle \boldsymbol{\xi}, \boldsymbol{w}\rangle) \mid \bar{\boldsymbol{A}}] \tag{25}$$

where, on the right-hand side in (25), the random vector $\boldsymbol{\xi}$ is seen as a fixed input to the conditional expectation given $\bar{\boldsymbol{A}}$. Or, to explain (25) more formally by using the expression (1) for the conditional expectation: (25) is equivalent to saying that, for all $\bar{\boldsymbol{A}}_{\boldsymbol{0}} \in \mathbb{Z}_q^{\bar{m}\times n}$ and $(\boldsymbol{s}_{\boldsymbol{0}}, \boldsymbol{e}_{\boldsymbol{0}}) \in [\pm\eta]^n \times [\pm\eta]^{\bar{m}}$ with $\mathbb{P}[\bar{\boldsymbol{A}} = \bar{\boldsymbol{A}}_{\boldsymbol{0}}, \boldsymbol{s} = \boldsymbol{s}_{\boldsymbol{0}}, \boldsymbol{e} = \boldsymbol{e}_{\boldsymbol{0}}] > 0$,

$$\mathbb{E}[\cos(2\pi\tfrac{1}{q}\langle \boldsymbol{\xi}, \boldsymbol{w}\rangle) \mid \bar{\boldsymbol{A}} = \bar{\boldsymbol{A}}_{\boldsymbol{0}},\ \boldsymbol{s} = \boldsymbol{s}_{\boldsymbol{0}},\ \boldsymbol{e} = \boldsymbol{e}_{\boldsymbol{0}}] \ = \ \mathbb{E}[\cos(2\pi\tfrac{1}{q}\langle \boldsymbol{\xi}_0, \boldsymbol{w}\rangle) \mid \bar{\boldsymbol{A}} = \bar{\boldsymbol{A}}_{\boldsymbol{0}}],$$

where $\boldsymbol{\xi}_0 \in [\pm\eta]^m$ denotes the value taken by $\boldsymbol{\xi}$ when $\boldsymbol{s} = \boldsymbol{s}_{\boldsymbol{0}}$ and $\boldsymbol{e} = \boldsymbol{e}_{\boldsymbol{0}}$.

Recall that, by Lemma 5, the conditional distribution of $\boldsymbol{w}$ given $\bar{\boldsymbol{A}}$ is $D_{\Lambda_t, qs}$. In particular $\boldsymbol{w} \in V$, so that $\langle \boldsymbol{\xi}, \boldsymbol{w}\rangle = \langle \pi(\boldsymbol{\xi}), \boldsymbol{w}\rangle$. Moreover, by (6), $\frac{1}{q}\Lambda_t$ is the dual lattice of $\pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$. Now we are in the position to apply the Poisson summation formula (Lemma 1). We obtain that

$$\begin{aligned}
\mathbb{E}[\cos(2\pi\tfrac{1}{q}\langle \boldsymbol{\xi}, \boldsymbol{w}\rangle) \mid \bar{\boldsymbol{A}}] &= \frac{\sum_{\boldsymbol{w}\in\Lambda_t} \cos(2\pi\frac{1}{q}\langle \pi(\boldsymbol{\xi}), \boldsymbol{w}\rangle)\rho_{qs}(\boldsymbol{w})}{\rho_{qs}(\Lambda_t)} = \frac{\sum_{\boldsymbol{w}\in\Lambda_t} \mathrm{e}^{2\pi\frac{1}{q}i\langle \pi(\boldsymbol{\xi}), \boldsymbol{w}\rangle}\rho_{qs}(\boldsymbol{w})}{\rho_{qs}(\Lambda_t)} \\
&= \frac{\sum_{\boldsymbol{v}\in\frac{1}{q}\Lambda_t} \mathrm{e}^{2\pi i\langle \pi(\boldsymbol{\xi}), \boldsymbol{v}\rangle}\rho_s(\boldsymbol{v})}{\rho_s(\frac{1}{q}\Lambda_t)} = \frac{\rho_{1/s}(\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})) - \pi(\boldsymbol{\xi}))}{\rho_{1/s}(\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})))} \\
&= f_{\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})),1/s}(-\pi(\boldsymbol{\xi})) = f_{\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})),1/s}(\pi(\boldsymbol{\xi})).
\end{aligned} \tag{26}$$

We then proceed as in [37, Corollary 1] to derive a simple lower bound. More precisely, we choose $\boldsymbol{y} \in \pi(L_q(\boldsymbol{A}_{\mathrm{dual}})) - \pi(\boldsymbol{\xi})$ with minimal norm. In particular,

we have that $\|\boldsymbol{y}\| = \mathrm{dist}(\pi(\boldsymbol{\xi}), \pi(L_q(\boldsymbol{A}_{\mathrm{dual}})))$. Then,

$$\begin{aligned}\rho_{1/s}(\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})) - \pi(\boldsymbol{\xi})) &= \rho_{1/s}(\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})) - \boldsymbol{y}) \\ &= \tfrac{1}{2}\textstyle\sum_{\boldsymbol{v}\in\pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))}(\rho_{1/s}(\boldsymbol{v}-\boldsymbol{y}) + \rho_{1/s}(-\boldsymbol{v}-\boldsymbol{y})) \\ &= \rho_{1/s}(\boldsymbol{y})\textstyle\sum_{\boldsymbol{v}\in\pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))}\rho_{1/s}(\boldsymbol{v})\cosh(2\pi s^2\langle\boldsymbol{v},\boldsymbol{y}\rangle) \\ &\geq \mathrm{e}^{-\pi s^2\|\boldsymbol{y}\|^2}\rho_{1/s}(\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})))\end{aligned} \tag{27}$$

Combining (24),(25), (26) and (27) concludes the proof.

8.2 Bounding the false-negative probability

Lemma 15. *Suppose the same setting as in Lemma 14. Then,*

$$\mathbb{P}[g_W(\boldsymbol{t}(\boldsymbol{s}_{\mathrm{guess}})) \leq c\cdot\mathcal{T}] \;\leq\; \mathbb{E}\big[\mathrm{e}^{-\frac{1}{2}N(1-c)^2\mathcal{T}^2}\big] + \tfrac{1}{2},$$

where the expectations are taken with respect to the randomness of N and $\mathcal{T}$, which both depend on t^* and thus on $\bar{\boldsymbol{A}}_{\mathrm{dual}}$.

Proof. As in Lemma 14, we abbreviate $\boldsymbol{\xi} := \big(\boldsymbol{e}, -\boldsymbol{s}_{\mathrm{dual}}\big)^{\mathsf{T}}$. We also, for the sake of readability, abbreviate $g_W := g_W(\boldsymbol{t}(\boldsymbol{s}_{\mathrm{guess}}))$. Then, using similar arguments as in [9, Theorem 4.9] or [35, Appendix B], we observe that via Lemma 14,

$$\begin{aligned}&\mathbb{P}[g_W \leq c\cdot\mathcal{T}] = \mathbb{P}\left[-g_W - \mathbb{E}[-g_W \,|\, (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] \;\geq\; \mathbb{E}[g_W \,|\, (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] - c\mathcal{T}\right] \\ &\leq \mathbb{P}\left[-g_W - \mathbb{E}[-g_W \,|\, (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] \;\geq\; \rho_{1/s}(\mathrm{dist}(\pi(\boldsymbol{\xi}), \pi(L_q(\boldsymbol{A}_{\mathrm{dual}})))) - c\mathcal{T}\right] \\ &\leq \mathbb{P}[-g_W - \mathbb{E}[-g_W \,|\, (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] \;\geq\; \rho_{1/s}(\|\pi(\boldsymbol{\xi})\|) - c\mathcal{T}],\end{aligned} \tag{28}$$

where we used the inequality $\mathrm{dist}(\pi(\boldsymbol{\xi}), \pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))) \leq \|\pi(\boldsymbol{\xi})\|$ in the last step. Then, the right-hand side of (28) is lower or equal to

$$\begin{aligned}&\mathbb{P}\big[-g_W - \mathbb{E}[-g_W \,|\, (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] \;\geq\; \rho_{1/s}(\|\pi(\boldsymbol{\xi})\|) - c\mathcal{T},\;\; \|\pi(\boldsymbol{\xi})\|^2 \leq \tfrac{t\eta}{2}(1+\gamma_t)\big] \\ &\qquad + \mathbb{P}\left[\|\pi(\boldsymbol{\xi})\|^2 > \tfrac{t\eta}{2}(1+\gamma_t)\right] \\ &\leq \mathbb{E}\big[\;\; \mathbb{P}\left[-g_W - \mathbb{E}[-g_W \,|\, (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})] \;\geq\; (1-c)\mathcal{T} \;\middle|\; (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})\right] \;\;\big] \\ &\qquad + \mathbb{P}\left[\|\pi(\boldsymbol{\xi})\|^2 > \tfrac{t\eta}{2}(1+\gamma_t)\right]\end{aligned}$$

Here we used that $\rho_{1/s}(\|\pi(\boldsymbol{\xi})\|) \geq \mathcal{T}$ whenever $\|\pi(\boldsymbol{\xi})\|^2 \leq \frac{t\eta}{2}(1+\gamma_t)$; cf. (10). Then, by applying Lemma 16 for $\mathbb{P}[\cdot \mid \bar{\boldsymbol{A}}_{\mathrm{dual}}]$ with $t = t^*$ and π (which are both functions of $\bar{\boldsymbol{A}}_{\mathrm{dual}}$) and $\boldsymbol{\xi}$ (which is independent of $\bar{\boldsymbol{A}}_{\mathrm{dual}}$ and distributed according to $\mathcal{C}(\eta)^{\otimes m}$), we infer that

$$\mathbb{P}\left[\|\pi(\boldsymbol{\xi})\|^2 > \tfrac{t\eta}{2}(1+\gamma_t)\right] = \mathbb{E}[\,\mathbb{P}\left[\|\pi(\boldsymbol{\xi})\|^2 > \tfrac{t\eta}{2}(1+\gamma_t) \;\middle|\; \bar{\boldsymbol{A}}_{\mathrm{dual}}\right]\,] \leq \tfrac{1}{2}$$

Finally, recalling Lemma 6, we can now apply Hoeffding's inequality (Lemma 21) for $S_N = -N\cdot g_W$ and for the conditional probability measure $\mathbb{P}[\cdot \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})]$, under which N and $\mathcal{T}$ are constants. This bounds the inner conditional probability by $\mathrm{e}^{-\frac{1}{2}N(1-c)^2\mathcal{T}^2}$ and concludes the proof.

8.3 Concentration of the projected error vector

Lemma 16. *Let $\boldsymbol{P} \in \mathbb{R}^{m\times m}$ be an orthogonal projection onto a subspace of dimension $t \in \{1,\dots,m\}$, and $\boldsymbol{\xi}$ be distributed according to the product distribution $\mathcal{C}(\eta)^{\otimes m}$, where we introduced $\mathcal{C}(\eta)$ in Definition 3 and m in (4). Then, $\mathbb{P}\left[\|\boldsymbol{P}\boldsymbol{\xi}\|^2 > \frac{t\eta}{2}(1+\gamma_t)\right] \le \frac{1}{2}$, where we defined $\gamma_t \in O(1/\sqrt{t})$ in (10).*

Proof. We write $\boldsymbol{\xi} = (\xi_i)_i$. By Definition 3, $\xi_i + \eta$ is binomially distributed with parameters 2η and $\frac{1}{2}$, so that $\mathbb{E}[\mathrm{e}^{\lambda\xi_i}] = \cosh(\lambda/2)^{2\eta} \le \mathrm{e}^{\eta\lambda^2/4}$ for all $\lambda \in \mathbb{R}$, where we used that $\cosh(x) \le \mathrm{e}^{x^2/2}$. By independence, $\mathbb{E}[\mathrm{e}^{\langle\boldsymbol{\alpha},\boldsymbol{\xi}\rangle}] \le \mathrm{e}^{\|\boldsymbol{\alpha}\|^2\eta/4}$ for all $\boldsymbol{\alpha} \in \mathbb{R}^m$. This is the hypothesis of [22, Theorem 2.1] with $\boldsymbol{\mu} = \boldsymbol{0}$ and $\sigma^2 = \frac{\eta}{2}$. Writing $\boldsymbol{P} = \boldsymbol{U}\boldsymbol{U}^\mathsf{T}$, where the columns of $\boldsymbol{U} \in \mathbb{R}^{m\times t}$ form an orthonormal basis of the subspace, we have $\boldsymbol{\Sigma} := \boldsymbol{P}^\mathsf{T}\boldsymbol{P} = \boldsymbol{P}^2 = \boldsymbol{P}$, so that $\mathrm{tr}(\boldsymbol{\Sigma}) = \mathrm{tr}(\boldsymbol{\Sigma}^2) = \mathrm{tr}(\boldsymbol{U}^\mathsf{T}\boldsymbol{U}) = t$ and $\|\boldsymbol{\Sigma}\| = 1$. Hence that theorem yields, for all $\tau > 0$,

$$\mathbb{P}\left[\|\boldsymbol{P}\boldsymbol{\xi}\|^2 > \tfrac{\eta}{2}\left(t + 2\sqrt{t\tau} + 2\tau\right)\right] \;\le\; \mathrm{e}^{-\tau}.$$

Choosing $\tau = \ln(2)$ concludes the proof.

9 The false-positive event

The only thing left to prove for Theorem 7 is the estimate on the probability of the false-positive event (FP). This is done in this section. We prove the bound in Section 9.1, and provide some auxiliary results in Section 9.2.

9.1 Bounding the false-positive probability

Lemma 17. *Suppose the same setting as in Theorem 7 with $t := t^*$, but again without (12) (cf. Lemma 14). Recall the definition of the rank function R given in the proof of Theorem 7. Then,*

$$\begin{aligned}\mathbb{P}[\,\exists\tilde{\boldsymbol{s}}_{\mathrm{guess}} \in [\pm\eta]^{n_{\mathrm{guess}}} : (R(\tilde{\boldsymbol{s}}_{\mathrm{guess}}) < R(\boldsymbol{s}_{\mathrm{guess}})) \wedge (g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\mathrm{guess}})) > c\mathcal{T})] \\ \le 2^{n_{\mathrm{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}\,\mathbb{E}\big[\mathrm{e}^{-\frac{1}{2}Nc^2\mathcal{T}^2}\big] + \tfrac{1}{m},\end{aligned} \tag{29}$$

where, as in Lemma 15, the expectation is taken with respect to the randomness of N and $\mathcal{T}$.

Proof. We abbreviate $\widetilde{g}_W := g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\mathrm{guess}}))$, $\tilde{R} := R(\tilde{\boldsymbol{s}}_{\mathrm{guess}})$ and $R := R(\boldsymbol{s}_{\mathrm{guess}})$, and we write $\tilde{f} := \mathbb{E}[\widetilde{g}_W \,|\, (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})]$. Note that the event $(\tilde{R} < R)$ is measurable with respect to the σ-algebra generated by $(\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})$, or, to put it in different terms, this event is totally determined if we know the realization value of the random vector $(\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})$. Then, by using a union bound and standard properties of the conditional expectation, the probability on the left-hand side of (29) is

upper bounded by

$$\begin{aligned}
&\textstyle\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}\in[\pm\eta]^{n_{\mathrm{guess}}}} \mathbb{P}\big[\,(\tilde{R}<R)\,\wedge\,(\widetilde{g}_W>c\mathcal{T})\big]\\
&=\textstyle\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}} \mathbb{E}\big[\,\mathbb{P}\big[\,(\tilde{R}<R)\,\wedge\,(\widetilde{g}_W>c\mathcal{T})\;\big|\;(\bar{\boldsymbol{A}},\boldsymbol{s},\boldsymbol{e})\,\big]\,\big]\\
&=\textstyle\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}} \mathbb{E}\big[\mathbb{1}_{(\tilde{R}<R)}\cdot\mathbb{P}\big[\;\widetilde{g}_W>c\mathcal{T}\;\big|\;(\bar{\boldsymbol{A}},\boldsymbol{s},\boldsymbol{e})\;\big]\;\big] \qquad (30)\\
&=\textstyle\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}} \mathbb{E}\big[\mathbb{1}_{(\tilde{R}<R)}\cdot\mathbb{1}_{(\tilde{f}<c\mathcal{T})}\cdot\mathbb{P}\big[\;\widetilde{g}_W-\tilde{f}>c\mathcal{T}-\tilde{f}\;\big|\;(\bar{\boldsymbol{A}},\boldsymbol{s},\boldsymbol{e})\;\big]\;\big]\\
&\quad+\textstyle\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}} \mathbb{E}\big[\mathbb{1}_{(\tilde{R}<R)}\cdot\mathbb{1}_{(\tilde{f}\geq c\mathcal{T})}\cdot\mathbb{P}\big[\;\widetilde{g}_W-\tilde{f}>c\mathcal{T}-\tilde{f}\;\big|\;(\bar{\boldsymbol{A}},\boldsymbol{s},\boldsymbol{e})\;\big]\;\big].
\end{aligned}$$

Recall Lemma 6. We apply Hoeffding's inequality (Lemma 21) to the inner conditional probability of the first term on the right-hand side of (30), with $S_N = N\cdot\widetilde{g}_W$ and for the conditional probability measure $\mathbb{P}[\,\cdot\mid(\bar{\boldsymbol{A}},\boldsymbol{s},\boldsymbol{e})]$, under which N, $\mathcal{T}$ and $\tilde{f}$ are constants. As $c\mathcal{T}-\tilde{f}>0$ on the event $(\tilde{f}<c\mathcal{T})$, this yields the bound $\mathrm{e}^{-\frac{1}{2}N(c\mathcal{T}-\tilde{f})^2}$. In the second term we bound the inner conditional probability trivially by 1. With ψ as in Lemma 22 for $a=c\mathcal{T}$, both bounds equal $\psi(\tilde{f})$ on the respective events. Hence, the left-hand side of (29) is at most $\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}}\mathbb{E}\big[\mathbb{1}_{(\tilde{R}<R)}\cdot\psi(\tilde{f})\big]$, which, in turn, by Lemma 22 is lower or equal to

$$\textstyle\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}} \mathbb{E}\big[\mathbb{1}_{(\tilde{R}<R)}\cdot\mathrm{e}^{-\frac{1}{2}Nc^2\mathcal{T}^2}\big]+\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}} \mathbb{E}\big[\mathbb{1}_{(\tilde{R}<R)}\cdot\tfrac{2}{c\mathcal{T}}\,\tilde{f}\big]. \qquad (31)$$

Here we have used that, by using the same arguments as in (24), (25) and (26),

$$\mathbb{E}\left[\widetilde{g}_W\,|\,(\bar{\boldsymbol{A}},\boldsymbol{s},\boldsymbol{e})\right] \;=\; f_{\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})),1/s}(\pi(\tilde{\boldsymbol{r}})) \;>\; 0 \qquad (32)$$

(as the periodic Gaussian function is positive), so that $\tilde{f}\geq 0$, which is required in Lemma 22. As $\boldsymbol{s}_{\mathrm{guess}}$ is independent of $\bar{\boldsymbol{A}}_{\mathrm{dual}}$, whereas N and $\mathcal{T}$ are functions of $\bar{\boldsymbol{A}}_{\mathrm{dual}}$, the first sum in (31) equals $\mathbb{E}\big[\mathrm{e}^{-\frac{1}{2}Nc^2\mathcal{T}^2}\big]\,(\mathbb{E}[R(\boldsymbol{s}_{\mathrm{guess}})]-1)$. The term $\mathbb{E}[R(\boldsymbol{s}_{\mathrm{guess}})]$ can be estimated by using Arikan's classical inequalities as in the proof of Theorem 9. The second sum is at most $\frac{1}{m}$ by Lemma 20, which concludes the proof.

9.2 Auxiliary results

Lemma 18 (Uniform distribution of the cosets). *Suppose the same setting as in Lemma 17. Let $\bar{\boldsymbol{A}}_{\boldsymbol{d}}\in\mathbb{Z}_q^{\bar{m}\times n_{\mathrm{dual}}}$, $\boldsymbol{s_0}\in[\pm\eta]^n$ and $\boldsymbol{e_0}\in[\pm\eta]^{\bar{m}}$ be values with $\mathbb{P}\left[\bar{\boldsymbol{A}}_{\mathrm{dual}}=\bar{\boldsymbol{A}}_{\boldsymbol{d}},\ \boldsymbol{s}=\boldsymbol{s_0},\ \boldsymbol{e}=\boldsymbol{e_0}\right]>0$, and let $\tilde{\boldsymbol{s}}_{\mathrm{guess}}\in[\pm\eta]^{n_{\mathrm{guess}}}$ be a fixed candidate with $\tilde{\boldsymbol{s}}_{\mathrm{guess}}\neq\boldsymbol{s_{0,g}}$, where $\boldsymbol{s_{0,g}}$ denotes the guessing part (i.e. the first n_{guess} coordinates) of the vector $\boldsymbol{s_0}$ (note that this is a condition on the parameters $\boldsymbol{s_0}$ and $\tilde{\boldsymbol{s}}_{\mathrm{guess}}$ of the lemma and not a further conditioning of the probability measure). Moreover, abbreviate*

$$\tilde{\boldsymbol{r}}:=\begin{pmatrix}\boldsymbol{e}+\bar{\boldsymbol{A}}_{\mathrm{guess}}(\boldsymbol{s}_{\mathrm{guess}}-\tilde{\boldsymbol{s}}_{\mathrm{guess}})\\ -\boldsymbol{s}_{\mathrm{dual}}\end{pmatrix}\in\mathbb{Z}^m.$$

Then, conditionally on the event $\{\bar{\boldsymbol{A}}_{\mathrm{dual}}=\bar{\boldsymbol{A}}_{\boldsymbol{d}},\boldsymbol{s}=\boldsymbol{s_0},\boldsymbol{e}=\boldsymbol{e_0}\}$, the coset $\tilde{\boldsymbol{r}}+L_q(\boldsymbol{A}_{\mathrm{dual}})$ is uniformly distributed on $\mathbb{Z}^m/L_q(\boldsymbol{A}_{\mathrm{dual}})$.

Proof. In the following, we abbreviate the probability measure $\mathbb{P}[\cdot \mid \bar{\boldsymbol{A}}_{\mathrm{dual}} = \bar{\boldsymbol{A}}_{\boldsymbol{d}}, \boldsymbol{s} = \boldsymbol{s_0}, \boldsymbol{e} = \boldsymbol{e_0}]$ as $\mathbb{P}^0$. Moreover, as in the proof of Lemma 14, let $\boldsymbol{\xi} := \big(\boldsymbol{e}, -\boldsymbol{s}_{\mathrm{dual}}\big)^{\mathsf{T}}$. For $\tilde{\boldsymbol{s}}_{\mathrm{guess}} \in [\pm\eta]^{n_{\mathrm{guess}}}$, we abbreviate $\tilde{\boldsymbol{\Delta}} := \boldsymbol{s}_{\mathrm{guess}} - \tilde{\boldsymbol{s}}_{\mathrm{guess}}$. Then, we have that $\tilde{\boldsymbol{r}} = \boldsymbol{\xi} + \boldsymbol{A}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}}$, where $\boldsymbol{A}_{\mathrm{guess}} := \left(\begin{smallmatrix} \bar{\boldsymbol{A}}_{\mathrm{guess}} \\ 0 \end{smallmatrix}\right) \in \mathbb{Z}_q^{m \times n_{\mathrm{guess}}}$.

Step 1: [Restriction to $\boldsymbol{A}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}} + L_q(\boldsymbol{A}_{\mathrm{dual}})$]
We can also write $\tilde{\boldsymbol{r}} + L_q(\boldsymbol{A}_{\mathrm{dual}}) = (\boldsymbol{\xi} + L_q(\boldsymbol{A}_{\mathrm{dual}})) + (\boldsymbol{A}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}} + L_q(\boldsymbol{A}_{\mathrm{dual}}))$. Note that (i) under $\mathbb{P}^0$, we have that $\boldsymbol{\xi} + L_q(\boldsymbol{A}_{\mathrm{dual}})$ is a fixed, non-random object, and, (ii) translation by the fixed element $\boldsymbol{\xi} + L_q(\boldsymbol{A}_{\mathrm{dual}})$ is a bijection of the finite group $\mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}})$. Combining (i) and (ii) shows that it suffices to prove that, under $\mathbb{P}^0$, the random coset $\boldsymbol{A}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}} + L_q(\boldsymbol{A}_{\mathrm{dual}})$ is uniformly distributed on $\mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}})$.

Step 2: [Preservation of the independence of the columns of $\bar{\boldsymbol{A}}_{\mathrm{guess}}$]
Using Definition 3, the matrix $\bar{\boldsymbol{A}}_{\mathrm{guess}}$ is independent of $(\bar{\boldsymbol{A}}_{\mathrm{dual}}, \boldsymbol{s}, \boldsymbol{e})$, so that, with respect to $\mathbb{P}^0$, the columns $(\bar{\boldsymbol{A}}_{\mathrm{guess}})_1, \ldots, (\bar{\boldsymbol{A}}_{\mathrm{guess}})_{n_{\mathrm{guess}}}$ are still independent and uniformly distributed on $\mathbb{Z}_q^{\bar{m}}$.

Step 3: [Uniform distribution of $\boldsymbol{A}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}}$]
Since $\tilde{\boldsymbol{\Delta}} = \boldsymbol{s_{0,g}} - \tilde{\boldsymbol{s}}_{\mathrm{guess}}$ has entries in $[\pm 2\eta]$ and is nonzero, and since $2\eta < q$ by Definition 3, we have $\tilde{\boldsymbol{\Delta}} \not\equiv \boldsymbol{0} \bmod q$. And as q is prime, there is an index k such that $\tilde{\Delta}_k$ is a unit in $\mathbb{Z}_q$. We aim to apply Lemma 25 with $G := \mathbb{Z}_q^{\bar{m}}$, $\boldsymbol{U} := (\bar{\boldsymbol{A}}_{\mathrm{guess}})_k$, $\boldsymbol{V} := \big((\bar{\boldsymbol{A}}_{\mathrm{guess}})_j\big)_{j \neq k}$, and $\psi(\boldsymbol{u}, \boldsymbol{v}) := \tilde{\Delta}_k\, \boldsymbol{u} + \sum_{j \neq k} \tilde{\Delta}_j\, \boldsymbol{v}_j$. Indeed, $\boldsymbol{U}$ is uniformly distributed on G and independent of $\boldsymbol{V}$. Moreover, for every fixed $\boldsymbol{v}$ the map $\psi(\cdot, \boldsymbol{v})$ is a bijection of G, because $\psi(\cdot, \boldsymbol{v})$ is the composition of the map $\boldsymbol{u} \mapsto \tilde{\Delta}_k \boldsymbol{u}$ (which is a bijection of G with inverse $\boldsymbol{u} \mapsto \tilde{\Delta}_k^{-1}\boldsymbol{u}$) with a translation of G. Since $\psi(\boldsymbol{U}, \boldsymbol{V}) = \bar{\boldsymbol{A}}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}}$, Lemma 25 shows that $\bar{\boldsymbol{A}}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}}$ is uniformly distributed on $\mathbb{Z}_q^{\bar{m}}$.

Step 4: [Uniform distribution of $\boldsymbol{A}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}} + L_q(\boldsymbol{A}_{\mathrm{dual}})$]
Combining Step 3 and Lemma 26, the coset $\boldsymbol{A}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}} + L_q(\boldsymbol{A}_{\mathrm{dual}}) = \Psi(\bar{\boldsymbol{A}}_{\mathrm{guess}}\tilde{\boldsymbol{\Delta}})$ (where we defined Ψ in Lemma 26) is uniformly distributed on $\mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}})$.

Lemma 19 (Conditional expectation of the periodic Gaussian function). *Suppose the same setting as in Lemma 18. Moreover, let Λ_t and V be defined as in* (3) *and* (5) *but with $\bar{\boldsymbol{A}}_{\mathrm{dual}} = \bar{\boldsymbol{A}}_{\boldsymbol{d}}$. Then,*

$$\begin{aligned} &\mathbb{E}\big[f_{\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})),1/s}(\pi(\tilde{\boldsymbol{r}})) \;\big|\; \bar{\boldsymbol{A}}_{\mathrm{dual}} = \bar{\boldsymbol{A}}_{\boldsymbol{d}},\, \boldsymbol{s} = \boldsymbol{s_0}, \boldsymbol{e} = \boldsymbol{e_0}\big] \\ &\qquad = \tfrac{\rho_s(\mathbb{Z}^m \cap V)}{\rho_{qs}(\Lambda_t)} \;\leq\; \rho_s(\mathbb{Z})^m\, \tfrac{\det(\Lambda_t)}{(qs)^t}. \end{aligned} \tag{33}$$

Proof. In the following, we abbreviate the probability measure $\mathbb{P}[\cdot \mid \bar{\boldsymbol{A}}_{\mathrm{dual}} = \bar{\boldsymbol{A}}_{\boldsymbol{d}}, \boldsymbol{s} = \boldsymbol{s_0}, \boldsymbol{e} = \boldsymbol{e_0}]$, under which Λ_t, V and π are fixed, as $\mathbb{P}^0$.

Step 1. [The cardinality of the quotient $G := \pi(\mathbb{Z}^m)/\pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$]
Note that, by (6), $\pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$ and $\pi(\mathbb{Z}^m)$ are lattices of rank t. Then, as $\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})) \subset \pi(\mathbb{Z}^m)$, $\pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$ is a sublattice of $\pi(\mathbb{Z}^m)$. Moreover, we observe that the map $\boldsymbol{x} + L_q(\boldsymbol{A}_{\mathrm{dual}}) \mapsto \pi(\boldsymbol{x}) + \pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$ is a surjective group homomorphism from $\mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}})$ onto G. Thus, as $|\mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}})| = q^{\bar{m}}$ by

Lemma 26, G is a finite group. We are now in the position to apply [28, Section 5] and infer that $|G| = \det(\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})))/\det(\pi(\mathbb{Z}^m))$.

Step 2. [Uniform distribution of $\pi(\tilde{\boldsymbol{r}}) + \pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$ under $\mathbb{P}^0$]
By Lemma 18, the random object $\tilde{\boldsymbol{r}} + L_q(\boldsymbol{A}_{\mathrm{dual}})$ is, under $\mathbb{P}^0$, uniformly distributed on $\mathbb{Z}^m/L_q(\boldsymbol{A}_{\mathrm{dual}})$. Therefore, using the same arguments as in [35, Lemma 10], we obtain that the coset $\pi(\tilde{\boldsymbol{r}}) + \pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$ is also uniformly distributed on G. Here we have again used that $\boldsymbol{x} + L_q(\boldsymbol{A}_{\mathrm{dual}}) \mapsto \pi(\boldsymbol{x}) + \pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$ is a surjective group homomorphism onto G.

Step 3. [Conclusion]
Combining Step 1 and Step 2 with the fact that the cosets of $\pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$ partition $\pi(\mathbb{Z}^m)$ (as $\pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$ is a subgroup of $\pi(\mathbb{Z}^m)$), the left-hand side of (33) equals

$$\frac{\frac{1}{|G|}\sum_{C\in G}\rho_{1/s}(C)}{\rho_{1/s}(\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})))} \;=\; \frac{\det(\pi(\mathbb{Z}^m))}{\det(\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})))}\cdot\frac{\rho_{1/s}(\pi(\mathbb{Z}^m))}{\rho_{1/s}(\pi(L_q(\boldsymbol{A}_{\mathrm{dual}})))} \;=\; \frac{\rho_s(\mathbb{Z}^m\cap V)}{\rho_{qs}(\Lambda_t)},$$

where, in the last step we applied Lemma 1 to $\pi(\mathbb{Z}^m)$ and to $\pi(L_q(\boldsymbol{A}_{\mathrm{dual}}))$, whose dual lattices are $\mathbb{Z}^m \cap V$ and $\frac{1}{q}\Lambda_t$ by (6), and we used that $\rho_s(\frac{1}{q}\Lambda_t) = \rho_{qs}(\Lambda_t)$. Finally, $\rho_s(\mathbb{Z}^m \cap V) \le \rho_s(\mathbb{Z}^m) = \rho_s(\mathbb{Z})^m$, and another application of Lemma 1 yields $\rho_{qs}(\Lambda_t) = \frac{(qs)^t}{\det(\Lambda_t)}\,\rho_{1/(qs)}(\widehat{\Lambda_t}) \ge \frac{(qs)^t}{\det(\Lambda_t)}$.

Lemma 20 (The remaining term from Lemma 17). *Suppose the same setting as in Lemma 17 and its proof. Then,*

$$\begin{aligned}\textstyle\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}\in[\pm\eta]^{n_{\mathrm{guess}}}} \mathbb{E}\Big[\mathbb{1}_{(R(\tilde{\boldsymbol{s}}_{\mathrm{guess}})<R(\boldsymbol{s}_{\mathrm{guess}}))}\cdot\tfrac{2}{c\mathcal{T}}\cdot\mathbb{E}[g_W(\boldsymbol{t}(\tilde{\boldsymbol{s}}_{\mathrm{guess}}))\,|\,(\bar{\boldsymbol{A}},\boldsymbol{s},\boldsymbol{e})]\Big] \\ \le\; \tfrac{1}{m\,2^{n_{\mathrm{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}}\left(\mathbb{E}\left[R(\boldsymbol{s}_{\mathrm{guess}})\right]-1\right) \;\le\; \tfrac{1}{m}.\end{aligned} \tag{34}$$

Proof. Recall the notation and setting of Lemma 18 and its proof. Moreover, let $\tilde{\boldsymbol{s}}_{\mathrm{guess}}$ be such that $R(\tilde{\boldsymbol{s}}_{\mathrm{guess}}) < R(\boldsymbol{s}_{\mathbf{0},\mathbf{g}})$. This implies in particular that $\tilde{\boldsymbol{s}}_{\mathrm{guess}} \neq \boldsymbol{s}_{\mathbf{0},\mathbf{g}}$. Finally, recall that $t = t^*$ is determined by $\bar{\boldsymbol{A}}_{\mathrm{dual}} = \bar{\boldsymbol{A}}_{\boldsymbol{d}}$, and that (11) holds for this t by the definition of t^*.

Let $\tilde{f} = \mathbb{E}[\tilde{g}_W\,|\,(\bar{\boldsymbol{A}},\boldsymbol{s},\boldsymbol{e})]$. Then, by (32), Lemma 19 and (11),

$$\tfrac{2}{c\mathcal{T}}\mathbb{E}[\tilde{f}\;|\;\bar{\boldsymbol{A}}_{\mathrm{dual}}=\bar{\boldsymbol{A}}_{\boldsymbol{d}},\,\boldsymbol{s}=\boldsymbol{s}_{\mathbf{0}},\boldsymbol{e}=\boldsymbol{e}_{\mathbf{0}}] \le \tfrac{2}{c\mathcal{T}}\rho_s(\mathbb{Z})^m\,\tfrac{\det(\Lambda_t)}{(qs)^t} \le \tfrac{1}{m\,2^{n_{\mathrm{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}}.$$

Using this inequality, abbreviating $\tilde{R} := R(\tilde{\boldsymbol{s}}_{\mathrm{guess}})$ and $R := R(\boldsymbol{s}_{\mathrm{guess}})$, and by conditioning on $(\bar{\boldsymbol{A}}_{\mathrm{dual}},\boldsymbol{s},\boldsymbol{e})$, the left-hand side of (34) is equal to

$$\begin{aligned}&\textstyle\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}}\mathbb{E}\big[\;\mathbb{E}[\,\mathbb{1}_{(\tilde{R}<R)}\cdot\tilde{f}\cdot\tfrac{2}{c\mathcal{T}}\;\big|(\bar{\boldsymbol{A}}_{\mathrm{dual}},\boldsymbol{s},\boldsymbol{e})]\;\big] \\ &\qquad=\textstyle\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}}\mathbb{E}\big[\mathbb{1}_{(\tilde{R}<R)}\cdot\mathbb{E}[\tilde{f}|(\bar{\boldsymbol{A}}_{\mathrm{dual}},\boldsymbol{s},\boldsymbol{e})]\cdot\tfrac{2}{c\mathcal{T}}\;\big] \le \tfrac{\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}}\mathbb{E}[\mathbb{1}_{(\tilde{R}<R)}]}{m\,2^{n_{\mathrm{guess}}\mathcal{H}_{1/2}(\mathcal{C}(\eta))}}\end{aligned}$$

where, similar to the proof of Lemma 17, we used that $(\tilde{R} < R)$ and $\mathcal{T}$ are measurable with respect to the σ-algebra generated by $(\bar{\boldsymbol{A}}_{\mathrm{dual}},\boldsymbol{s},\boldsymbol{e})$. Finally, we note that $\sum_{\tilde{\boldsymbol{s}}_{\mathrm{guess}}}\mathbb{E}[\mathbb{1}_{(\tilde{R}<R)}] = \mathbb{E}[R(\boldsymbol{s}_{\mathrm{guess}})]-1$, and that we can bound this term using Arikan's inequality (as in the proof of Theorem 9).

A Hoeffding's concentration inequality

Lemma 21 (Hoeffding's inequality). *Let $(X_i)_{i=1}^N$ be independent random variables such that $a_i \le X_i \le b_i$ almost surely for some $a_i, b_i \in \mathbb{R}$ for all i. Let $S_N = \sum_i X_i$. Then, for all $r > 0$, $\mathbb{P}\left[S_N - \mathbb{E}\left[S_N\right] \ge r\right] \;\le\; \mathrm{e}^{-2r^2\left(\sum_{i=1}^N (b_i - a_i)^2\right)^{-1}}$.*

Note that, in Section 8 and Section 9, we applied this inequality not for $\mathbb{P}[\cdot]$ but for the conditional probability measure $\mathbb{P}[\cdot \mid (\bar{\boldsymbol{A}}, \boldsymbol{s}, \boldsymbol{e})]$. This is admissible as the proof of Hoeffding's inequality only relies on elementary properties of the expectation and the probability measure, which also hold true for their conditional versions.

In the proof of Lemma 17 we combine Hoeffding's inequality with the following elementary inequality.

Lemma 22. *Let $a, N > 0$, and let $\psi(x) := \mathrm{e}^{-\frac{1}{2}N((a-x)_+)^2}$ for $x \ge 0$, where $y_+ := \max\{y, 0\}$. Then $\psi(x) \le \mathrm{e}^{-\frac{1}{2}Na^2} + \frac{2x}{a}$ for all $x \ge 0$.*

Proof. If $x \ge \frac{a}{2}$, then the right-hand side is at least $1 \ge \psi(x)$, and for $x = 0$ there is nothing to prove. If $0 < x < \frac{a}{2}$, then, by the mean value theorem, $\psi(x) - \psi(0) = x\,\psi'(\xi)$ for some $\xi \in (0, x)$, where $\psi'(\xi) = Nu\,\mathrm{e}^{-\frac{1}{2}Nu^2}$ with $u := a - \xi > \frac{a}{2}$. As $y\,\mathrm{e}^{-y/2} \le \frac{2}{\mathrm{e}}$ for all $y \ge 0$ (we apply this with $y = Nu^2$), we obtain $\psi'(\xi) \le \frac{2}{\mathrm{e}\,u} < \frac{4}{\mathrm{e}\,a} < \frac{2}{a}$. As $\psi(0) = \mathrm{e}^{-\frac{1}{2}Na^2}$, this concludes the proof.

B Some conditional expectation calculus

Lemma 23. *Let $\boldsymbol{\Gamma}$ and $\boldsymbol{\Theta}$ be finite random variables, let $\boldsymbol{X}$ be an arbitrary random variable, and suppose that $\boldsymbol{X}$ and $\boldsymbol{\Theta}$ are independent with respect to $\mathbb{P}[\,\cdot \mid \boldsymbol{\Gamma}]$. Then, for every bounded measurable function G,*

$$\mathbb{E}\left[G(\boldsymbol{X}, \boldsymbol{\Gamma}, \boldsymbol{\Theta}) \mid (\boldsymbol{\Gamma}, \boldsymbol{\Theta})\right](\cdot) = g\big(\boldsymbol{\Gamma}(\cdot), \boldsymbol{\Theta}(\cdot)\big),$$

where $g(l, \theta) := \mathbb{E}\left[G(\boldsymbol{X}, l, \theta) \mid \boldsymbol{\Gamma} = l\right]$.

Proof. By (1) it suffices to fix values l of $\boldsymbol{\Gamma}$ and θ of $\boldsymbol{\Theta}$ with $\mathbb{P}[A] > 0$, where $A := \{\boldsymbol{\Gamma} = l, \boldsymbol{\Theta} = \theta\}$, and to show that $\mathbb{E}[G(\boldsymbol{X}, \boldsymbol{\Gamma}, \boldsymbol{\Theta}) \mid A] = g(l, \theta)$. We have that

$$\mathbb{E}\left[G(\boldsymbol{X}, \boldsymbol{\Gamma}, \boldsymbol{\Theta}) \mid A\right] \;=\; \tfrac{\mathbb{E}[G(\boldsymbol{X}, \boldsymbol{\Gamma}, \boldsymbol{\Theta})\mathbb{1}_A]}{\mathbb{P}[A]} \;=\; \tfrac{\mathbb{E}[G(\boldsymbol{X}, l, \theta)\mathbb{1}_A]}{\mathbb{P}[A]} \;=\; \mathbb{E}\left[G(\boldsymbol{X}, l, \theta) \mid A\right].$$

It remains to show that the right-hand side is equal to $g(l, \theta)$. Applying the assumed independence under $\mathbb{P}[\,\cdot \mid \boldsymbol{\Gamma} = l]$ yields

$$\mathbb{E}\left[G(\boldsymbol{X}, l, \theta)\mathbb{1}_{\{\boldsymbol{\Theta} = \theta\}} \mid \boldsymbol{\Gamma} = l\right] \;=\; \mathbb{E}\left[G(\boldsymbol{X}, l, \theta) \mid \boldsymbol{\Gamma} = l\right] \cdot \mathbb{P}\left[\boldsymbol{\Theta} = \theta \mid \boldsymbol{\Gamma} = l\right],$$

which, after dividing by $\mathbb{P}[\boldsymbol{\Theta} = \theta \mid \boldsymbol{\Gamma} = l] > 0$, shows $\mathbb{E}[G(\boldsymbol{X}, l, \theta) \mid A] = g(l, \theta)$.

Corollary 24. *In the setting of Lemma 23, let $\boldsymbol{X} = (\boldsymbol{X}_1, \dots, \boldsymbol{X}_N)$, where N is allowed to be a function of $\boldsymbol{\Gamma}$, and assume in addition that $\boldsymbol{X}_1, \dots, \boldsymbol{X}_N$ are independent with respect to $\mathbb{P}[\cdot \mid \boldsymbol{\Gamma}]$. Then, for all bounded measurable $G_1, \dots, G_N$,*

$$\mathbb{E}\big[\textstyle\prod_{j=1}^{N} G_j(\boldsymbol{X}_j, \boldsymbol{\Gamma}, \boldsymbol{\Theta}) \mid (\boldsymbol{\Gamma}, \boldsymbol{\Theta})\big] \;=\; \prod_{j=1}^{N} \mathbb{E}\big[G_j(\boldsymbol{X}_j, \boldsymbol{\Gamma}, \boldsymbol{\Theta}) \mid (\boldsymbol{\Gamma}, \boldsymbol{\Theta})\big].$$

In particular, for arbitrary bounded, measurable functions $f_1, \dots, f_N$, we have that the family $\{f_j(\boldsymbol{X}_j, \boldsymbol{\Gamma}, \boldsymbol{\Theta})\}_{j=1}^{N}$ is independent with respect to $\mathbb{P}\big[\cdot \mid (\boldsymbol{\Gamma}, \boldsymbol{\Theta})\big]$.

Proof. Let l, θ, A be as in the proof of Lemma 23. Using the independence of $\boldsymbol{X}_1, \dots, \boldsymbol{X}_N$ under $\mathbb{P}[\cdot \mid \boldsymbol{\Gamma} = l]$ yields that

$$\mathbb{E}[\textstyle\prod_j G_j(\boldsymbol{X}_j, l, \theta) \mid \boldsymbol{\Gamma} = l] \;=\; \prod_j \mathbb{E}[G_j(\boldsymbol{X}_j, l, \theta) \mid \boldsymbol{\Gamma} = l].$$

Applying Lemma 23 to $G = \prod_j G_j$ for the left-hand side, and to $G = G_j$ for each j for the right-hand side, implies that

$$\mathbb{E}\big[\textstyle\prod_j G_j(\boldsymbol{X}_j, \boldsymbol{\Gamma}, \boldsymbol{\Theta}) \mid A\big] \;=\; \prod_j \mathbb{E}\big[G_j(\boldsymbol{X}_j, \boldsymbol{\Gamma}, \boldsymbol{\Theta}) \mid A\big].$$

which concludes the proof of the first claim. For the second claim, we apply the first one with $G_j = h_j \circ f_j$ for arbitrary bounded, measurable functions $h_1, \dots, h_N$. This leads exactly to the defining property of independence of a family of random objects. In other words, $h_1, \dots, h_N$ act as the arbitrary test functions for proving the independence claim.

C Preserving uniformity under bijective functions

Lemma 25. *Let G be a finite non-empty set, let S be a countable set, let $\boldsymbol{U}$ be a random variable that is uniformly distributed on G, and let $\boldsymbol{V}$ be a random variable with values in S. Suppose that $\boldsymbol{V}$ is independent of $\boldsymbol{U}$. Let $\psi : G \times S \to G$ be a map such that, for every $v \in S$, the map $\psi(\cdot, v) : G \to G$ is a bijection. Then $\psi(\boldsymbol{U}, \boldsymbol{V})$ is uniformly distributed on G and independent of $\boldsymbol{V}$.*

Proof. Let $y \in G$ and let $v \in S$ with $\mathbb{P}[\boldsymbol{V} = v] > 0$. Since $\psi(\cdot, v)$ is a bijection of G, there is a unique $u_{y,v} \in G$ with $\psi(u_{y,v}, v) = y$, and therefore

$$\begin{aligned} \mathbb{P}\left[\psi(\boldsymbol{U}, \boldsymbol{V}) = y,\ \boldsymbol{V} = v\right] \;&=\; \mathbb{P}\left[\psi(\boldsymbol{U}, v) = y,\ \boldsymbol{V} = v\right] \\ = \mathbb{P}\left[\boldsymbol{U} = u_{y,v},\ \boldsymbol{V} = v\right] &= \mathbb{P}\left[\boldsymbol{U} = u_{y,v}\right] \cdot \mathbb{P}\left[\boldsymbol{V} = v\right] = \tfrac{1}{|G|} \cdot \mathbb{P}\left[\boldsymbol{V} = v\right], \end{aligned} \tag{35}$$

where we used the independence of $\boldsymbol{U}$ and $\boldsymbol{V}$ in the second-last step, and the uniformity of $\boldsymbol{U}$ in the last step. Summing (35) over all $v \in S$ with $\mathbb{P}[\boldsymbol{V} = v] > 0$ yields $\mathbb{P}[\psi(\boldsymbol{U}, \boldsymbol{V}) = y] = \frac{1}{|G|}$, so that $\psi(\boldsymbol{U}, \boldsymbol{V})$ is uniformly distributed on G. Substituting this back into (35) gives $\mathbb{P}[\psi(\boldsymbol{U}, \boldsymbol{V}) = y, \boldsymbol{V} = v] = \mathbb{P}[\psi(\boldsymbol{U}, \boldsymbol{V}) = y] \cdot \mathbb{P}[\boldsymbol{V} = v]$ for all y and v, which is the asserted independence.

D On the number of cosets

Lemma 26. *The map* $\Psi : \mathbb{Z}_q^{\bar{m}} \longrightarrow \mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}})$, $\Psi(\boldsymbol{u}) = \binom{\boldsymbol{u}}{\boldsymbol{0}} + L_q(\boldsymbol{A}_{\mathrm{dual}})$, *is a group isomorphism. In particular,* $|\mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}})| = q^{\bar{m}}$.

Proof. Ψ is clearly a group homomorphism. Moreover, it is well-defined, since $q\mathbb{Z}^m \subseteq L_q(\boldsymbol{A}_{\mathrm{dual}})$. Now suppose that $\boldsymbol{u}$ is in its kernel, i.e. $(\boldsymbol{u}, \boldsymbol{0}) \in L_q(\boldsymbol{A}_{\mathrm{dual}})$. This implies that $(\boldsymbol{u}, \boldsymbol{0}) = \boldsymbol{A}_{\mathrm{dual}}\boldsymbol{x} \bmod q$ for some $\boldsymbol{x} \in \mathbb{Z}^{n_{\mathrm{dual}}}$. And since, by (3), the lower block of $\boldsymbol{A}_{\mathrm{dual}}$ equals $\boldsymbol{I}_{n_{\mathrm{dual}}}$, we have that $\boldsymbol{x} \equiv \boldsymbol{0} \bmod q$ and thus $\boldsymbol{u} \equiv \boldsymbol{0} \bmod q$. We obtain the injectivity of Ψ.

It remains to show that $|\mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}})| = q^{\bar{m}} = |\mathbb{Z}_q^{\bar{m}}|$, which implies the surjectivity of Ψ. Using the same arguments as above, we obtain that the map $\boldsymbol{x} \mapsto \boldsymbol{A}_{\mathrm{dual}}\boldsymbol{x} \bmod q$ is injective on $\mathbb{Z}_q^{n_{\mathrm{dual}}}$. This implies that the set $J := \{\boldsymbol{A}_{\mathrm{dual}}\boldsymbol{x} \bmod q \mid \boldsymbol{x} \in \mathbb{Z}_q^{n_{\mathrm{dual}}}\}$, a subgroup of $\mathbb{Z}_q^m$, has cardinality $|J| = q^{n_{\mathrm{dual}}}$. Next, we define $\varphi : \mathbb{Z}^m \to \mathbb{Z}_q^m / J$ by $\varphi(\boldsymbol{y}) = \boldsymbol{y} \bmod q + J$. Then, φ is obviously a surjective homomorphism. Moreover, $\boldsymbol{y} \in \ker\varphi$ if and only if $\boldsymbol{y} \bmod q \in J$. The latter can easily be seen to be equivalent to saying that $\boldsymbol{y} \in L_q(\boldsymbol{A}_{\mathrm{dual}})$. Therefore, $\ker\varphi = L_q(\boldsymbol{A}_{\mathrm{dual}})$. We thus, by the first isomorphism theorem, obtain that $\mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}}) = \mathbb{Z}^m / \ker\varphi \cong \mathrm{im}\varphi = \mathbb{Z}_q^m / J$. This, in turn, implies via Lagrange's theorem that

$$|\mathbb{Z}^m / L_q(\boldsymbol{A}_{\mathrm{dual}})| = \left|\mathbb{Z}_q^m / J\right| = \frac{|\mathbb{Z}_q^m|}{|J|} = \frac{q^m}{q^{n_{\mathrm{dual}}}} = q^{\bar{m}} = |\mathbb{Z}_q^{\bar{m}}|,$$

and thus concludes the proof.

AI Disclosure

An LLM was used in preparing this paper to improve English wording, for proofreading, and for generating the code for the parameter searches and for the experiments in Section 7. In addition, the proofs of some of the lemmas were prepared with LLM assistance. The main ideas and strategies of this paper were designed by the author, who takes full responsibility for its entire content.

References

1. Aharonov, D., Regev, O.: Lattice problems in NP ∩ coNP. J. ACM **52**(5), 749–765 (2005). https://doi.org/10.1145/1089023.1089025
2. Albrecht, M.R., Gheorghiu, V., Postlethwaite, E.W., Schanck, J.M.: Estimating quantum speedups for lattice sieves. In: Moriai, S., Wang, H. (eds.) Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part II. Lecture Notes in Computer Science, vol. 12492, pp. 583–613. Springer (2020). https://doi.org/10.1007/978-3-030-64834-3_20, https://doi.org/10.1007/978-3-030-64834-3_20
3. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. J. Math. Cryptol. **9**(3), 169–203 (2015), http://www.degruyter.com/view/j/jmc.2015.9.issue-3/jmc-2015-0016/jmc-2015-0016.xml

4. Arbel, J., Marchal, O., Nguyen, H.D.: On strict sub-Gaussianity, optimal proxy variance and symmetry for bounded random variables. ESAIM: Probability and Statistics **24**, 39–55 (2020). https://doi.org/10.1051/ps/2019018
5. Arikan, E.: An inequality on guessing and its application to sequential decoding. IEEE Transactions on Information Theory **42**(1), 99–105 (1996)
6. Arora, S., Ge, R.: New algorithms for learning in presence of errors. In: International Colloquium on Automata, Languages, and Programming. pp. 403–415. Springer (2011)
7. Bai, S., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS-Dilithium: Algorithm specifications and supporting documentation (version 3.1). Submission to the NIST Post-Quantum Cryptography Standardization Process (2021), https://pq-crystals.org/dilithium/data/dilithium-specification-round3-20210208.pdf
8. Bashiri, K., Glaser, T., May, A., Nowakowski, J.: Super-quadratic quantum speed-ups and guessing many likely keys. In: Daemen, J., Thomé, E. (eds.) Advances in Cryptology - EUROCRYPT 2026 - 45th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Rome, Italy, May 10-14, 2026, Proceedings, Part I. Lecture Notes in Computer Science, vol. 16541, pp. 457–480. Springer (2026). https://doi.org/10.1007/978-3-032-25291-3_16
9. Bashiri, K., Wiemers, A.: On the independence heuristic in the dual attack. J. Math. Cryptol. **19**(1) (2025). https://doi.org/10.1515/JMC-2024-0028
10. Becker, A., Ducas, L., Gama, N., Laarhoven, T.: New directions in nearest neighbor searching with applications to lattice sieving. In: Krauthgamer, R. (ed.) Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016. pp. 10–24. SIAM (2016). https://doi.org/10.1137/1.9781611974331.CH2, https://doi.org/10.1137/1.9781611974331.ch2
11. Billingsley, P.: Probability and Measure. John Wiley & Sons (2012)
12. de Boer, K., van Woerden, W.P.J.: Lattice-based cryptography: A survey on the security of the lattice-based NIST finalists. IACR Cryptol. ePrint Arch. **2025**, 304 (2025), https://eprint.iacr.org/2025/304
13. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors (2013). https://doi.org/10.1145/2488608.2488680
14. 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
15. Carrier, K., Meyer-Hilfiger, C., Shen, Y., Tillich, J.: Assessing the impact of a variant of MATZOV's dual attack on Kyber. In: Kalai, Y.T., Kamara, S.F. (eds.) Advances in Cryptology - CRYPTO 2025 - 45th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2025, Proceedings, Part I. Lecture Notes in Computer Science, vol. 16000, pp. 444–476. Springer (2025). https://doi.org/10.1007/978-3-032-01855-7_15
16. 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 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part III. Lecture Notes in Computer Science, vol. 14083, pp. 37–69. Springer (2023). https://doi.org/10.1007/978-3-031-38548-3_2
17. Ducas, L., Pulles, L.N.: Accurate score prediction for dual-sieve attacks. J. Cryptol. **39**(1), 8 (2026). https://doi.org/10.1007/S00145-025-09560-7

18. Gama, N., Nguyen, P.Q.: Predicting lattice reduction. In: Smart, N.P. (ed.) Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings. Lecture Notes in Computer Science, vol. 4965, pp. 31–51. Springer (2008). https://doi.org/10.1007/978-3-540-78967-3_3
19. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Dwork, C. (ed.) Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008. pp. 197–206. ACM (2008). https://doi.org/10.1145/1374376.1374407
20. 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 - 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6-10, 2021, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 13093, pp. 33–62. Springer (2021). https://doi.org/10.1007/978-3-030-92068-5_2
21. Hou, J., Jiang, H., Ogilvie, T.: Careful with the ring! Concrete hardness gaps between LWE and MLWE. In: Heninger, N., Rosulek, M. (eds.) Advances in Cryptology - CRYPTO 2026 - 46th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-20, 2026, Proceedings, Part III. Lecture Notes in Computer Science, vol. 16802, pp. 462–497. Springer (2026). https://doi.org/10.1007/978-3-032-35377-1_15
22. Hsu, D., Kakade, S.M., Zhang, T.: A tail inequality for quadratic forms of subgaussian random vectors. Electronic Communications in Probability **17**(52), 1–6 (2012). https://doi.org/10.1214/ECP.v17-2079
23. Karenin, A., Kirshanova, E., Nowakowski, J., May, A.: Fast slicer for Batch-CVP: Making lattice hybrid attacks practical. In: Hanaoka, G., Yang, B. (eds.) Advances in Cryptology - ASIACRYPT 2025 - 31st International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, VIC, Australia, December 8-12, 2025, Proceedings, Part III. Lecture Notes in Computer Science, vol. 16247, pp. 100–132. Springer (2025). https://doi.org/10.1007/978-981-95-5099-9_4
24. Kirchner, P., Fouque, P.: An improved BKW algorithm for LWE with applications to cryptography and lattices. In: Gennaro, R., Robshaw, M. (eds.) Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I. Lecture Notes in Computer Science, vol. 9215, pp. 43–62. Springer (2015). https://doi.org/10.1007/978-3-662-47989-6_3
25. Klein, P.N.: Finding the closest lattice vector when it's unusually close. In: Shmoys, D.B. (ed.) Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, January 9-11, 2000, San Francisco, CA, USA. pp. 937–941. ACM/SIAM (2000), http://dl.acm.org/citation.cfm?id=338219.338661
26. Laarhoven, T., Walter, M.: Dual lattice attacks for closest vector problems (with preprocessing). In: Paterson, K.G. (ed.) Topics in Cryptology - CT-RSA 2021 - Cryptographers' Track at the RSA Conference 2021, Virtual Event, May 17-20, 2021, Proceedings. Lecture Notes in Computer Science, vol. 12704, pp. 478–502. Springer (2021). https://doi.org/10.1007/978-3-030-75539-3_20
27. Lenstra, A.K., Lenstra, Jr., H.W., Lovász, L.: Factoring polynomials with rational coefficients. Mathematische Annalen **261**(4), 515–534 (1982). https://doi.org/10.1007/BF01457454

28. Lenstra, Jr., H.W.: Lattices. In: Buhler, J.P., Stevenhagen, P. (eds.) Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography, Mathematical Sciences Research Institute Publications, vol. 44, pp. 127–181. Cambridge University Press (2008)
29. Li, J., Wang, L.P., Wang, H.: Provable FFT-accelerated dual attack on LWE using the general dual lattice. Cryptology ePrint Archive, Paper 2026/2186 (2026), https://eprint.iacr.org/2026/2186
30. Lyubashevsky, V.: Basic Lattice Cryptography: The concepts behind Kyber (ML-KEM) and Dilithium (ML-DSA). Cryptology ePrint Archive, Paper 2024/1287 (2024), https://eprint.iacr.org/2024/1287
31. Menezes, A.: A gentle introduction to lattice-based cryptography. Cryptology ePrint Archive, Paper 2026/1098 (2026), https://eprint.iacr.org/2026/1098
32. National Institute of Standards and Technology: Module-Lattice-Based Digital Signature Standard. Federal Information Processing Standards Publication (FIPS) NIST FIPS 204, U.S. Department of Commerce, Washington, D.C. (2024). https://doi.org/10.6028/NIST.FIPS.204
33. National Institute of Standards and Technology: Module-Lattice-Based Key-Encapsulation Mechanism Standard. Federal Information Processing Standards Publication (FIPS) NIST FIPS 203, U.S. Department of Commerce, Washington, D.C. (2024). https://doi.org/10.6028/NIST.FIPS.203
34. Peikert, C.: A decade of lattice cryptography. Found. Trends Theor. Comput. Sci. **10**(4), 283–424 (2016). https://doi.org/10.1561/0400000074
35. Pouly, A., Shen, Y.: Provable dual attacks on learning with errors. In: Joye, M., Leander, G. (eds.) Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part VI. Lecture Notes in Computer Science, vol. 14656, pp. 256–285. Springer (2024). https://doi.org/10.1007/978-3-031-58754-2_10
36. Qu, H., Tian, C., Wang, G., Xu, G.: Provable dual attack on LWE via lattice projection. Cryptology ePrint Archive, Paper 2026/2117 (2026), https://eprint.iacr.org/2026/2117
37. Qu, H., Xu, G.: On the provable dual attack for LWE by modulus switching. In: Hanaoka, G., Yang, B. (eds.) Advances in Cryptology - ASIACRYPT 2025 - 31st International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, VIC, Australia, December 8-12, 2025, Proceedings, Part III. Lecture Notes in Computer Science, vol. 16247, pp. 34–64. Springer (2025). https://doi.org/10.1007/978-981-95-5099-9_2
38. Serre, J.P.: A Course in Arithmetic, Graduate Texts in Mathematics, vol. 7. Springer-Verlag, New York (1973)
39. Wang, R., Wang, Z., Zheng, Q., Zhao, X.: LaMS: A p-adic layered modulus switching for provable dual attacks on LWE. IACR Cryptol. ePrint Arch. **2026**, 1326 (2026), https://eprint.iacr.org/2026/1326
40. Wang, Z., Ling, C.: Lattice Gaussian sampling by Markov chain Monte Carlo: Bounded distance decoding and trapdoor sampling. IEEE Trans. Inf. Theory **65**(6), 3630–3645 (2019). https://doi.org/10.1109/TIT.2019.2901497