





Provable Dual Attacks on Learning with Errors

Amaury Pouly¹  and Yixin Shen² 

¹ Centre National de la Recherche Scientifique (CNRS), Paris, France
amaury.pouly@cnrs.fr

² King's College London, London, UK
yixin.shen@kcl.ac.uk

Abstract. Learning with Errors (LWE) is an important problem for post-quantum cryptography (PQC) that underlines the security of several NIST PQC selected algorithms. Several recent papers [7, 25], [16, 32] have claimed improvements on the complexity of so-called dual attacks on LWE. These improvements make dual attacks comparable to or even better than primal attacks in certain parameter regimes. Unfortunately, those improvements rely on a number of untested and hard-to-test statistical assumptions. Furthermore, a recent paper [20] claims that the whole premise of those improvements might be incorrect.

The goal of this paper is to improve the situation by proving the correctness of a dual attack without relying on any statistical assumption. Although our attack is greatly simplified compared to the recent ones, it shares many important technical elements with those attacks and can serve as a basis for the analysis of more advanced attacks. We provide some rough estimates on the complexity of our simplified attack on Kyber using a Monte Carlo Markov Chain discrete Gaussian sampler.

Our main contribution is to clearly identify a set of parameters under which our attack (and presumably other recent dual attacks) can work. Furthermore, our analysis completely departs from the existing statistics-based analysis and is instead rooted in geometry. We also compare the regime in which our algorithm works to the “contradictory regime” of [20]. We observe that those two regimes are essentially complementary.

Finally, we give a quantum version of our algorithm to speed up the computation. The algorithm is inspired by [10] but is completely formal and does not rely on any heuristics.

Keywords: Learning with Errors · Dual attack · Lattice-based cryptography · Quantum algorithm

1 Introduction

The Learning With Errors (LWE) problem [40] has become central to the security of several cryptosystems. Most notably, Kyber (public-key encryption) and Dilithium (signature) have been selected by the NIST for the Post-Quantum Cryptography (PQC) Standardization and rely on algebraic version of LWE for

their security proofs. Other advanced cryptographic primitives such as FHE can be built with LWE [15]. This makes LWE security estimates critical for the future of PQC. The *search LWE problem* asks to recover the secret \mathbf{s} given (\mathbf{A}, \mathbf{b}) where $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$, \mathbf{A} is a matrix chosen uniformly at random and \mathbf{e} has small entries (more details in Sect. 2.1).

There are two main approaches to attack the LWE problem: so-called primal and dual attacks. In this paper, we will exclusively focus on dual attacks which have recently attracted some interest due to significant claimed improvements in their complexity. Both primal and dual attacks rely on the BKZ lattice reduction algorithm [43] to obtain short vectors in lattices. The fundamental idea of dual attacks is to use short vectors in the dual of the lattice to detect whether points are close to the lattice or not, an idea that can be traced back to [5]. This allows us to solve the *distinguishing LWE problem* where one is asked to detect whether a sample comes from an LWE distribution, or a uniform distribution [35]. In conjunction with some guessing step, this allows one to recover part of the secret by trying several values until we get a point close to the lattice. By repeating this operation a couple of times, we can solve the search LWE problem.

Originally, the main limiting factor (on the complexity) of dual attacks was the need to compute one short vector (a very expensive operation) for every few LWE samples (more details in Sect. 3) and compute a score for each secret guess. Since then, a series of improvements have found their way into these attacks. First, a series of works on lattice sieving have shown [13, 36, 38] that those algorithms produce not only one but in fact exponentially many short vectors “for free”. [11] suggested that this idea could be used in dual attacks but it appears that [23] was the first paper to try to analyze it. Independently, [7] used a “re-randomization” technique to produce many short vectors from a single BKZ reduced basis. All those techniques claim to reduce the complexity of attacks although the correctness relies on an unproven assumption about the quality of those many short vectors. Then [25] noted that instead of computing the score for each secret guess separately, all the scores can be computed at once using a discrete Fourier transform (DFT), essentially reducing the cost to that of a single guess. Following this work, a technical report by the MATZOV group [32] has claimed further improvements by the use of a “modulus switching” technique¹ that significantly reduces the size of the DFT. Two recent work have modified this attack to include a quantum [10] and lattice coding [16] speed up.

One issue with the papers above is that the number of statistical assumptions that are necessary to justify the correctness of the algorithms has grown significantly, notably in [32]. While certain assumptions could probably be justified (almost) formally, others are subject to more controversy [20]. In particular, the most controversial aspect of [25, 32] is that the attack only uses a few LWE samples and that all the (exponentially-many) short vectors are derived from those samples which therefore are not statistically independent. When using a small number of LWE samples, the problem becomes very close to the *Bounded Dis-*

¹ A modulus switching technique was also suggested in [25] but it is unclear to us how it compares to [32], and [20] suggests that they are different.

tance Decoding which has been extensively studied. The status of [7] is unclear because it computes exponentially many short vectors from exponentially many samples, but the ratio of the number of short vectors to the number of samples is also exponential so the issue of the statistical independence remains but it does not seem as problematic. This makes it unclear whether an argument like that of [20] applies to such a case.

The purpose of this paper is to encourage a more rigorous analysis of dual attacks on LWE to better understand under what set of parameters they provably work. We note in that regard that a recently accepted paper at TCC 2023 [33] has focused on similar problems in statistical decoding/“dual attacks” in coding theory. The authors claim in the conclusion that at least part of their results apply to lattice dual attack. We believe that it would indeed be interesting to see what this approach yields for lattices, however we point out that the notion of dual attack that the authors have in mind looks quite different from the one in this paper. In short, and with our notations, the “dual attack” of [33] would be akin to splitting \mathbf{A} horizontally instead of vertically. This splitting would not correspond anymore to a decomposition of $L_q(\mathbf{A})$ as $L_q(\mathbf{A}_{\text{guess}}) + L_q(\mathbf{A}_{\text{dual}})$ and therefore looks incompatible with existing works on dual attacks on LWE. Furthermore, our understanding of [33] is that generating parity check vectors \mathbf{h} corresponds to generating many short dual vectors in $L_q^\perp(\mathbf{A})$, independently of the splitting of \mathbf{A} . This is completely at odds with lattice dual attacks where we split \mathbf{A} to generate dual vectors in $L_q^\perp(\mathbf{A}_{\text{dual}})$ which is much cheaper. Overall it looks like [33] might be a completely different kind of dual attack. See [39, Appendix A] for more details.

1.1 Contributions

The main contribution of this paper is to provide a *completely formal, non-asymptotic analysis* of a simplified dual attack. To simplify the presentation, we do not include elements such as the guessing complexity and modulus switching² to focus on the most controversial element, namely the fact that the attack only uses m LWE samples (with m not much bigger than the dimension n of the samples) and that all the short vectors are derived from those m samples.

Our approach completely departs from the existing statistics-based attacks and is instead rooted in geometry. This allows us to obtain a relatively short proof and leverage existing results on the geometry of lattices.

One of the most important technical contribution of this paper is to make completely clear (Theorem 5) under what choice of parameters the attack works, without any statistical assumption. As far as we are aware, no other dual attack has been formally analyzed in this way. We believe that this is important since virtually all algorithms in the literature rely on statistical assumptions that clearly cannot hold for all parameter regimes but without a proper analysis, it is impossible to tell when and why they hold.

² See Sect. 7 for more details about modulus switching.

We also provide some new results on random q -ary lattices in a similar spirit to that of Siegel, Rogers and Macbeath [31, 41, 44]. This allows us to obtain some sharper bounds on λ_1 for random q -ary lattices and show that the Gaussian Heuristic is quite tight for such lattices. This heuristic is usually considered valid for “random” lattices and has been extensively tested. Up to our knowledge, the only formal analysis of λ_1 for random q -ary lattice is in [47, Lemma 7.9.2] which only analyzes the expected value and therefore provides a much weaker bound on λ_1 . We refer to Sect. 2.3 for more details.

Finally, we give a quantum version of our algorithm to speed up the computation. The algorithm is inspired by [10] and reuses some technical lemmas to speed up the computation of sums of cosines that appear in the algorithm. Similarly to our classical algorithm, we prove that our quantum algorithm is correct without relying on any heuristics.

1.2 Comparison with [20]’s Contradictory Regime

A recent paper [20] has claimed that virtually all recent dual attacks rely on an incorrect statistical assumption and that they are, therefore, probably incorrect. They do so by formalizing what they claim to be the key statistical assumption of those paper, and show that for the parameter regime of the attacks, it falls into what they call the “contradictory regime”, a regime where this assumption can be proven not to hold.

As a byproduct of our analysis, we are able to compare the regime in which our analysis works with the contradictory regime of [20]. Interestingly, the two are essentially complementary with a small gap inbetween. This suggests that our analysis and that of [20] are quite tight and provide an almost complete characterization of when dual attacks work in our simplified setting. However, we nuance this conclusion by noting that the statistical model used in [20] to argue about the contradiction does not seem to match what happens in our algorithm. We refer to Sect. 6 for more details.

1.3 Organisation of the Paper

In Sect. 2, we introduce the various technical elements that are necessary to analyse the dual attack. In Sect. 3, we first present a basic dual attack whose purpose is to introduce the reader to the ideas of dual attacks without overwhelming them with technical details. This dual attack is very naive and computes one short vector per LWE sample, in the spirit of [5]. We emphasize that this attack and Theorem 4 are not new but that our analysis is significantly simpler than in previous papers. In Sect. 4, we introduce our simplified dual attack in the spirit of [32] and formally analyse its correctness without assumption. We provide some rough estimates on the complexity of our attack on Kyber using a Monte Carlo Markov Chain discrete Gaussian sampler. In Sect. 5, we give a quantum version of the algorithm from Sect. 4 and prove its correctness. In Sect. 6, we compare our regime with that of [20]. Finally, in Sect. 7, we describe what we believe is the main obstacle to develop a formal analysis of the full algorithm in [32].

2 Preliminaries

We denote vectors and matrices in bold case. We denote by \mathbf{x}^T the transpose of the (column) vector \mathbf{x} , which is therefore a row vector. We denote by \mathbf{I}_n the identity matrix of size $n \times n$. For any vector $\mathbf{x} \in \mathbb{R}^n$, we denote by $\|\mathbf{x}\|$ its Euclidean norm. We denote by $\langle \mathbf{x}, \mathbf{y} \rangle$ the scalar product between two vectors \mathbf{x} and \mathbf{y} . For any function $f : \mathbb{R}^n \rightarrow \mathbb{C}$, we denote by \widehat{f} its Fourier transform over \mathbb{R}^n defined by $\widehat{f}(\mathbf{x}) = \int_{\mathbb{R}^n} f(\mathbf{y}) e^{-2i\pi \langle \mathbf{x}, \mathbf{y} \rangle} d\mathbf{x}$. For any $n \in \mathbb{N}$ and $R > 0$, we denote by $B_n(R)$ (resp. $\overline{B}_n(R)$) the open (resp. closed) ball of radius R in \mathbb{R}^n . We also let $B_n^{\mathbb{Z}}(R) = B_n(R) \cap \mathbb{Z}^n$ be the set of integers points in this ball, and similarly for $\overline{B}_n^{\mathbb{Z}}(R)$. For any two distributions P and Q , we denote by $d_{\text{TV}}(P, Q)$ the statistical distance (or total variation distance) between P and Q . For any finite set X , we denote by $\mathcal{U}(X)$ the uniform distribution over X .

2.1 LWE

Let $n, m, q \in \mathbb{N}$ and let χ_e be a distribution over \mathbb{Z}_q , which we call the *noise distribution*. For every vector $\mathbf{s} \in \mathbb{Z}_q^n$, we denote by $\text{LWE}(m, \mathbf{s}, \chi_e)$ the probability distribution on $\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ obtained by sampling a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ uniformly at random, sampling a vector $\mathbf{e} \in \mathbb{Z}_q^m$ according to χ_e^m , and outputting (\mathbf{A}, \mathbf{b}) where $\mathbf{b} := \mathbf{A}\mathbf{s} + \mathbf{e}$. This is the “matrix form” for the LWE distribution where each pair (\mathbf{A}, \mathbf{b}) encodes m LWE samples $\mathbf{b}_i = \langle \mathbf{A}_i, \mathbf{s} \rangle + \mathbf{e}_i$ in the sense of [40]. We have chosen this formalism because it is simpler for dual attacks. The value of m is typically in the order of n and depends on the cryptosystem.

The search LWE problem is to find \mathbf{s} given oracle access to a sampler for $\text{LWE}(m, \mathbf{s}, \chi_e)$. The decision LWE problem is to decide, given oracle access to either $\text{LWE}(m, \mathbf{s}, \chi_e)$ or $\mathcal{U}(\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m)$, which one it is. In practical scenarios, the attacker may not have access to the sampler but rather only possess a limited number LWE samples. In this case, the search LWE problem asks, given those LWE samples, to recover \mathbf{s} if possible.

The LWE secret \mathbf{s} is usually generated according to a distribution χ_s over \mathbb{Z}_q^n . One can therefore, in principle, analyse the success probability of an algorithm for search/decision LWE on a distribution $\text{LWE}(m, \mathbf{s}, \chi_e)$ where $\mathbf{s} \leftarrow \chi_s^n$. In this paper, we will not need to make any assumption on the distribution of the secret since our algorithms work for every secret.

2.2 Discrete Gaussian Distribution

Let $n \in \mathbb{N}$ and $s > 0$. For any $\mathbf{x} \in \mathbb{R}^n$, we let $\rho_s(\mathbf{x}) := e^{-\pi \|\mathbf{x}\|^2 / s^2}$. As usual, we extend to ρ_s to sets by $\rho_s(X) = \sum_{\mathbf{x} \in X} \rho_s(\mathbf{x})$ for any set X . For any lattice $L \subset \mathbb{R}^n$, we denote the *discrete Gaussian distribution* over L by $D_{L,s}(\mathbf{x}) = \frac{\rho_s(\mathbf{x})}{\rho_s(L)}$ for any $\mathbf{x} \in L$. We denote $D_{L,1}$ by D_L for simplicity.

In general, the smaller s is, the harder it is to construct a sampler for $D_{L,s}$. The notion of smoothing parameter [34] captures the idea that sampling for a value of s above this threshold is significantly easier than sampling below because the distribution looks more like a continuous Gaussian. There are many algorithms to sample above the smoothing parameter [14, 24, 28], including a

time-space trade-off [3]. Sampling below the smoothing parameter is much more challenging and usually inefficient [4]. At the extreme, sampling for sufficiently small values of s allows one to solve the Shortest Vector problem (SVP) [4] which is known to be NP-hard under randomized reduction [6]. The Monte Carlo Markov Chain based algorithm of [46] works for all values of s but the complexity significantly depends on s . We will use this algorithm in this paper.

Theorem 1 ([46, Theorem 1, (8), (23) and (24)³]). *There is an algorithm that given a basis \mathbf{B} of a lattice $L \subset \mathbb{R}^n$, any $\varepsilon > 0$ and any $s > 0$, returns a sample according to some distribution $\mathcal{D}_{L,s,\varepsilon}$ such that $d_{\text{TV}}(\mathcal{D}_{L,s,\varepsilon}, D_{L,s}) \leq \varepsilon$. This algorithm runs in time $\ln(\frac{1}{\varepsilon}) \cdot \frac{1}{\Delta} \cdot \text{poly}(n)$ where $\frac{1}{\Delta} = \frac{\prod_{i=1}^n \rho_{s/\|\tilde{\mathbf{b}}_i\|(\mathbb{Z})}}{\rho_s(L)}$ and $\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_n$ are the Gram-Schmidt vectors of \mathbf{B} .*

For any $q \in \mathbb{N}$, we denote by $D_{\mathbb{Z}_q^n,s}$ the modular discrete Gaussian distribution over \mathbb{Z}_q^n defined by $D_{\mathbb{Z}_q^n,s}(\mathbf{x}) = \frac{\rho_s(\mathbf{x}+q\mathbb{Z}^n)}{\rho_s(\mathbb{Z}^n)}$ for any $\mathbf{x} \in \mathbb{Z}_q^n$. We define the periodic Gaussian function $f_{L,s} : \mathbb{R}^n \rightarrow \mathbb{R}$ by $f_{L,s}(\mathbf{t}) = \frac{\rho_s(L+\mathbf{t})}{\rho_s(L)}$. We have $f_{L/s,1}(\mathbf{t}/s) = f_{L,s}(\mathbf{t})$. In the following, we denote $f_{L,1}$ as f_L .

Lemma 1 ([17, Lemma 2.14]). *For any $L, s > 0, \mathbf{x} \in \mathbb{R}^n, f_{L,s}(\mathbf{x}) \geq \rho_s(\mathbf{x})$.*

Lemma 2 ([12, Lemma 7], see also [45, Theorem 1.3.4]). *For any lattice $L \subset \mathbb{R}^n, \mathbf{x} \in \mathbb{R}^n$ and $u \geq 1/\sqrt{2\pi}, \rho_s((L - \mathbf{x}) \setminus B_n(us\sqrt{n})) \leq (u\sqrt{2\pi}ee^{-\pi u^2})^n \rho_s(L)$.*

Corollary 1 ([45, Corollary 1.3.5]). *For any lattice $L \subset \mathbb{R}^n, \mathbf{t} \in \mathbb{R}^n$ and $r \geq \delta := s\sqrt{n}/2\pi, \rho_s((L - \mathbf{t}) \setminus B_n(r)) \leq \rho_s(r - \delta)\rho_s(L)$.*

Lemma 3 ([5, Claim 4.1]). *For any lattice L and $s > 0$, we have $\widehat{f_{L,s}} = D_{\widehat{L},1/s}$ which is a probability measure over the dual lattice \widehat{L} .*

2.3 Lattices

We denote by $\widehat{L} = \{\mathbf{x} \in \text{span}(L) : \forall \mathbf{y} \in L, \langle \mathbf{y}, \mathbf{x} \rangle \in \mathbb{Z}\}$ the dual of a lattice $L \subset \mathbb{R}^n$. We denote by $L^* = L \setminus \{\mathbf{0}\}$ the set of nonzero vectors of a lattice L . We denote by $\lambda_1(L)$ the length a shortest nonzero vector in L .

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

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

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

We refer the reader to [22], [47, Section 2.5.1] or [35] for more details on those constructions. Note that, equivalently, we can write $L_q(\mathbf{A}) = \mathbf{A}\mathbb{Z}_q^k + q\mathbb{Z}^n$. It is well-know that for any q -ary lattice L , there exists \mathbf{A} and \mathbf{B} such that $L = L_q(\mathbf{A}) =$

³ [46] uses the normal distribution $e^{-\|\mathbf{x}\|^2/2\sigma^2}$ so $s = \sqrt{2\pi}\sigma$ with our notations.

$L_q^\perp(\mathbf{B})$, and that $\widehat{L_q^\perp(\mathbf{A})} = \frac{1}{q}L_q(\mathbf{A})$. Furthermore $\det(L_q(\mathbf{A})) = q^{n-\text{rk } \mathbf{A}} \geq q^{n-k}$ and therefore $\det(L_q^\perp(\mathbf{A})) = q^{\text{rk } \mathbf{A}} \leq q^k$. Finally, since \mathbb{Z}_q is a field, a random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times k}$ has full rank (equal to k) with probability at least $1 - kq^{k-1-n}$. We will consider the distributions $\mathcal{L}_{n,k,q}$ and $\mathcal{L}_{n,k,q}^\perp$ of q -ary lattices defined over the set of integer lattices by

$$\begin{aligned} \mathcal{L}_{n,k,q}(L) &= \Pr_{\mathbf{A} \leftarrow \mathfrak{U}(\mathbb{Z}_q^{n \times k})} [L = L_q(\mathbf{A})], \\ \mathcal{L}_{n,k,q}^\perp(L) &= \Pr_{\mathbf{A} \leftarrow \mathfrak{U}(\mathbb{Z}_q^{n \times (n-k)})} [L = L_q^\perp(\mathbf{A})]. \end{aligned}$$

In other words, the distribution is obtained by taking a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times k}$ with uniform and i.i.d entries, and looking at the q -ary lattice generated by \mathbf{A} ; and similarly for the orthogonal version. Note that contrary to the Loeliger ensemble $\mathbb{L}_{n,k,q,1}$, we do not have the rescaling factor $q^{1-k/n}$, see e.g. [47, Definition 7.9.2]. It will be more convenient to use $\mathcal{L}_{n,k,q}^\perp$ for proofs, but we often want to apply them for $\mathcal{L}_{n,k,q}$. Whenever neither k nor $n - k$ are too small, those two distributions are very close. The following lemma was inspired by [19, Lemma 2] which does not contain any proof.

Lemma 4 ([39, Appendix C.1]). *Let $n \in \mathbb{N}$, $1 \leq k \leq n$ and q be a prime power. Then $d_{\text{TV}}(\mathcal{L}_{n,k,q}^\perp, \mathcal{L}_{n,k,q}) \leq \text{poly}(n, k) q^{-\min(k, n-k)}$.*

Those distributions satisfy good uniformity properties when q goes to infinity. In particular, the following theorem shows that we can compute statistical properties of lattices sampled according to $\mathcal{L}_{n,k,q}^\perp$. The first part of this theorem is close to [30, Theorem 1]. This result is in some sense the q -ary version of the result by Siegel on random (real) lattices and its generalization by Rogers and Macbeath [31, 41, 44].

Theorem 2 ([39, Appendix C.2]). *Let $n \in \mathbb{N}$, $1 \leq k \leq n$ and q be a prime power. Let $1 \leq p \leq n$ and $f : (\mathbb{Z}_q^n)^p \rightarrow \mathbb{R}$, then*

$$\mathbb{E}_{L \leftarrow \mathfrak{L}_{n,k,q}^\perp} \left[\sum_{\mathbf{x}_1, \dots, \mathbf{x}_p \in L} f(\mathbf{x}_1, \dots, \mathbf{x}_p) \right] = \sum_{\mathbf{x}_1, \dots, \mathbf{x}_p \in \mathbb{Z}^n} q^{(k-n)r(\mathbf{x}_1, \dots, \mathbf{x}_p)} f(\mathbf{x}_1, \dots, \mathbf{x}_p)$$

where $r(\mathbf{x}_1, \dots, \mathbf{x}_p) := \text{rk}_{\mathbb{Z}_q^n}(\mathbf{x}_1, \dots, \mathbf{x}_p)$ is the rank of the $\mathbf{x}_i \bmod q$ over \mathbb{Z}_q^n .

We can apply this theorem to bound the expected number of lattice points in a ball, and therefore obtain bounds on λ_1 .

Theorem 3 ([39, Appendix C.3]). *Let $n \in \mathbb{N}$, $1 \leq k \leq n$ and q be a prime power. For any $0 < r \leq q$,*

$$\begin{aligned} \mathbb{E}_{L \leftarrow \mathfrak{L}_{n,k,q}^\perp} [|L^* \cap B_n(r)|] &= q^{k-n} (|B_n^\mathbb{Z}(r)| - 1), \\ \mathbb{V}_{L \leftarrow \mathfrak{L}_{n,k,q}^\perp} [|L^* \cap B_n(r)|] &\leq q^{k-n} (q - 1) (|B_n^\mathbb{Z}(r)| - 1). \end{aligned}$$

In particular, if $|B_n^\mathbb{Z}(r)| \leq q^{n-k}$, then $\Pr_{L \leftarrow \mathfrak{L}_{n,k,q}^\perp} [\lambda_1(L) \leq r] \leq q^{1+k-n} |B_n^\mathbb{Z}(r)|$.

Recall that the *Gaussian heuristic* says that for a “random” lattice L , $\lambda_1(L)$ is approximately

$$\text{GH}(L) := \left(\frac{\text{vol}(B_n)}{\det(L)} \right)^{-1/n} = \frac{\det(L)^{1/n} \Gamma(1 + \frac{n}{2})^{1/n}}{\sqrt{\pi}} \approx \det(L)^{1/n} \sqrt{\frac{n}{2\pi e}}.$$

This heuristic is usually considered valid for “random” lattices and has been extensively tested. Up to our knowledge, the only formal analysis of λ_1 for random q -ary lattice is in [47, Lemma 7.9.2] which only analyzes the expected value and not the variance. The following corollary shows that this heuristic is indeed very sharp for random q -ary lattices.

Corollary 2 (Informal, [39, Appendix C.4]). *Let $n \in \mathbb{N}$, $1 \leq k \leq n$ and q be a prime power. Let $\alpha \in [0, 1]$ and $r = q^{1-k/n} \text{vol}(B_n)^{-1/n}$. Under the assumption that $|B_n^{\mathbb{Z}}(\alpha r)| \approx \text{vol}(B_n(\alpha r))$, which holds when $\alpha r \gg \sqrt{n}$, we have*

$$\Pr_{L \leftarrow \mathcal{L}_{n,k,q}} [\lambda_1(L) \leq \alpha \text{GH}(L)] \lesssim q\alpha^n.$$

Lemma 5 (The Pointwise Approximation Lemma [5, Lemma 1.3], modified). *Let $L \subset \mathbb{R}^n$ be a lattice, and $h : \mathbb{R}^n \rightarrow \mathbb{R}$ a L -periodic function whose Fourier series \hat{h} is a probability measure over \hat{L} . Let $N \in \mathbb{N}$, $\delta > 0$ and $X \subseteq \mathbb{R}^n$ a finite set. Let $W = (\mathbf{w}_1, \dots, \mathbf{w}_N)$ be a list of vectors in the dual lattice chosen randomly and independently from the distribution \hat{h} . Then with probability at least $1 - |X|2^{-\Omega(N\delta^2)}$, $h_W(\mathbf{x}) := \frac{1}{N} \sum_{i=1}^N \cos(2\pi \langle \mathbf{w}_i, \mathbf{x} \rangle)$ satisfies that $|h_W(\mathbf{x}) - h(\mathbf{x})| \leq \delta$ for all $\mathbf{x} \in L + X$.*

Proof. The proof is the one in [5] with the following modifications. Let $\delta > 0$. For any $\mathbf{x} \in \mathbb{R}^n$, Hoeffding’s inequality guarantees that the mean of N samples is not within a window of δ of the correct expectation with probability at most $2^{-\Omega(N\delta^2)}$. Since f is periodic over the lattice L , it suffices to check that the inequality that we want holds for all $\mathbf{x} \in X$. Hence, by a union bound, the probability that the approximation is within a window δ of the correct expectation for all $\mathbf{x} \in X$ simultaneously is at least $1 - |X|2^{-\Omega(N\delta^2)}$. \square

2.4 Short Vector Sampling

For the purpose of this paper, we will only need to know that there is a way to sample relatively short vectors (SV) in a lattice and we will treat such an algorithm as a black box. Since such an algorithm would typically be parametrized (see below), we introduce an integer parameter β to capture this fact.

Black Box 1. *For any integers $n \leq m$, β and prime power q , there exists a deterministic algorithm \mathcal{B} and two functions T_{SV} and ℓ_{SV} such that when \mathcal{B} is given $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, it returns a nonzero vector in $L_q^\perp(\mathbf{A})$ in time $T_{\text{SV}}(m, \beta, q^n)$ and $\mathbb{E}_{\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}} [\|\mathcal{B}(\mathbf{A})\|^2] \leq \ell_{\text{SV}}(m, \beta, q^n)^2$.*

One way to implement this black box is to use lattice reduction algorithms such as BKZ: they provide a very flexible way to take a basis of lattice and compute relatively short vectors in this lattice. Since the literature on this topic is quite extensive and there are many cost models associated to that task, we refer the reader to e.g. [25] for more details. For simplicity, we assume that the algorithm is deterministic but we could make it probabilistic by adding random coins to the input of the algorithm and take those into account in the expected value. In the case of BKZ, the parameter β is the block size.

3 Basic Dual Attack

In this section, we present a basic dual attack whose purpose is to introduce the reader to the ideas of dual attacks without overwhelming them with technical details. This dual attack is very naive and assumes that we access to essentially an unlimited number of samples. It computes one short vector per m LWE samples, in the spirit of [5]. We emphasize that this attack and Theorem 4 are not new but that our analysis is significantly simpler than in previous papers.

Fix $\mathbf{s} \in \mathbb{Z}_q^n$ an unknown secret and (\mathbf{A}, \mathbf{b}) some LWE samples. Recall that $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ for some unknown $\mathbf{e} \in \mathbb{Z}_q^m$. We split the secret \mathbf{s} into two parts $\mathbf{s}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ and $\mathbf{s}_{\text{dual}} \in \mathbb{Z}_q^{n_{\text{dual}}}$ where $n = n_{\text{guess}} + n_{\text{dual}}$. The matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is correspondingly split into two parts:

$$\mathbf{A} = [\mathbf{A}_{\text{guess}} \ \mathbf{A}_{\text{dual}}], \quad \mathbf{s} = \begin{bmatrix} \mathbf{s}_{\text{guess}} \\ \mathbf{s}_{\text{dual}} \end{bmatrix}. \tag{1}$$

Therefore, $\mathbf{b} = \mathbf{A}_{\text{guess}}\mathbf{s}_{\text{guess}} + \mathbf{A}_{\text{dual}}\mathbf{s}_{\text{dual}} + \mathbf{e}$. The algorithm now makes a guess $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ on the value of $\mathbf{s}_{\text{guess}}$ and tries to check whether this guess is correct. Consider the lattice

$$L_q^\perp(\mathbf{A}_{\text{dual}}) = \{ \mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \mathbf{A}_{\text{dual}} = \mathbf{0} \pmod q \}. \tag{2}$$

By the inequalities of Sect. 2.3, we have that $\det(L_q^\perp(\mathbf{A}_{\text{dual}})) \leq q^{n_{\text{dual}}}$. Check that for any $\mathbf{x} \in L_q^\perp(\mathbf{A}_{\text{dual}})$,

$$\mathbf{x}^T \mathbf{b} = \mathbf{x}^T \mathbf{A}_{\text{guess}}\mathbf{s}_{\text{guess}} + \mathbf{x}^T \mathbf{A}_{\text{dual}}\mathbf{s}_{\text{dual}} + \mathbf{x}^T \mathbf{e} = \mathbf{x}^T \mathbf{A}_{\text{guess}}\mathbf{s}_{\text{guess}} + \mathbf{x}^T \mathbf{e} \pmod q.$$

Therefore, $\mathbf{x}^T(\mathbf{b} - \mathbf{A}_{\text{guess}}\tilde{\mathbf{s}}_{\text{guess}}) = \mathbf{x}^T \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}}) + \mathbf{x}^T \mathbf{e} \pmod q$. The main observation is now that:

- if the guess is correct ($\tilde{\mathbf{s}}_{\text{guess}} = \mathbf{s}_{\text{guess}}$) then $\mathbf{x}^T(\mathbf{b} - \mathbf{A}_{\text{guess}}\tilde{\mathbf{s}}_{\text{guess}}) = \mathbf{x}^T \mathbf{e} \pmod q$ follows roughly a modular Gaussian distribution,
- if the guess is incorrect ($\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$) then it follows a uniform distribution because $\mathbf{x} \neq \mathbf{0}$ and \mathbf{A} was chosen uniformly at random.

A crucial ingredient in the reasoning above is the length of \mathbf{x} . Indeed, the scalar product $\mathbf{x}^T \mathbf{e}$ will follow a modular Gaussian whose deviation is proportional to $\|\mathbf{x}\|$. This is where the BKZ lattice reduction algorithm usually comes in: from a basis of $L_q^\perp(\mathbf{A}_{\text{dual}})$, we compute a short vector \mathbf{x} using Black box 1.

The algorithm for this attack is described in Algorithm 1. We group many LWE samples in N tuples of m samples which we write in matrix form. We then compute one dual vector for each tuple of m LWE samples as explained above. In this attack, the value of m can be chosen arbitrarily and there usually is an optimal value of m that can be computed based on the complexity of computing a short vector, *i.e.* it depends on the specific instantiation of Black box 1.

While this kind of attack is already known to be correct, we reprove it for several reasons. First, we are not satisfied with the informal treatment of the proof in the literature. Second, our proof does not use any assumption whereas most papers in the literature use the Central Limit Theorem or approximate sums of Gaussian as a Gaussian at some point (see [39, Section 2.4]). Figure 1 gives a high level view of the variable involved and their dependencies.

Theorem 4 ([39, Appendix B]). *Let n, m, β be integers, q be a prime power, $n_{\text{guess}} + n_{\text{dual}} = n$, $\mathbf{s} \in \mathbb{Z}_q^n$, $\sigma_e > 0$ and $N \in \mathbb{N}$. Let $0 < \delta < \varepsilon$ where $\varepsilon := \exp(-\pi\sigma_e^2 \ell_{\text{SV}}(m, \beta, q^{n_{\text{dual}}})^2/q^2)$ and ℓ_{SV} comes from Black box 1. Let $(\mathbf{A}^{(1)}, \mathbf{b}^{(1)}), \dots, (\mathbf{A}^{(N)}, \mathbf{b}^{(N)})$ be samples from $\text{LWE}(m, \mathbf{s}, D_{\mathbb{Z}_q, \sigma_e})$, then Algorithm 1 on $(m, n_{\text{guess}}, n_{\text{dual}}, q, \delta, N, (\mathbf{A}^{(i)}, \mathbf{b}^{(i)})_i)$ runs in time $\text{poly}(m, n) \cdot (N \cdot T_{\text{SV}}(m, \beta, q^n) + q^{n_{\text{guess}}})$ and returns $\mathbf{s}_{\text{guess}}$ with probability at least $1 - \exp\left(-\frac{N(\varepsilon - \delta)^2}{2}\right) - (q^{n_{\text{guess}}} - 1) \exp\left(-\frac{N\delta^2}{2}\right)$ over the choice of the $(\mathbf{A}^{(i)}, \mathbf{b}^{(i)})$.*

Remark 1. As expected, we recover the well-known fact that for the attack to succeed with constant probability, we can take $\delta = \varepsilon/2$ and then we need at least $N = \frac{8n_{\text{guess}} \log(q) + \Omega(1)}{\varepsilon^2}$ samples. Furthermore, a careful look at the proof shows that Black box 1 can be weakened even further to only require an inequality on the moment-generating function of $\|\mathcal{B}(\mathbf{A})\|^2$.

Algorithm 1: Basic dual attack

Input: $m, n = n_{\text{guess}} + n_{\text{dual}}$ (see (1)), q prime power, $\delta > 0$ and $N \in \mathbb{N}$.
Input: list of N LWE samples $(\mathbf{A}^{(1)}, \mathbf{b}^{(1)}), \dots, (\mathbf{A}^{(N)}, \mathbf{b}^{(N)})$.
Output: (Guess of) the first n_{guess} coordinates of the secret or \perp .

- 1 **for** j from 1 to N **do**
- 2 Compute a basis of $L_q^\perp(\mathbf{A}_{\text{dual}}^{(j)})$;
- 3 Compute a short vector $\mathbf{x}_j \in L_q^\perp(\mathbf{A}_{\text{dual}}^{(j)})$ using Black box 1;
- 4 **for** $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ **do**
- 5 Compute the list y_1, \dots, y_N where $y_j = \mathbf{x}_j^T(\mathbf{b}^{(j)} - \mathbf{A}_{\text{guess}}^{(j)} \tilde{\mathbf{s}}_{\text{guess}})$;
- 6 **if** $\frac{1}{N} \sum_{j=1}^N \cos(2\pi y_j/q) \geq \delta$ **then return** $\tilde{\mathbf{s}}_{\text{guess}}$;
- 7 **return** \perp

4 Modern Dual Attack

The main limitation of the basic dual attack is the requirement to compute one short vector for each tuple of m LWE samples. Looking at Fig. 1, this is necessary to ensure the statistical independence of the variables that go into

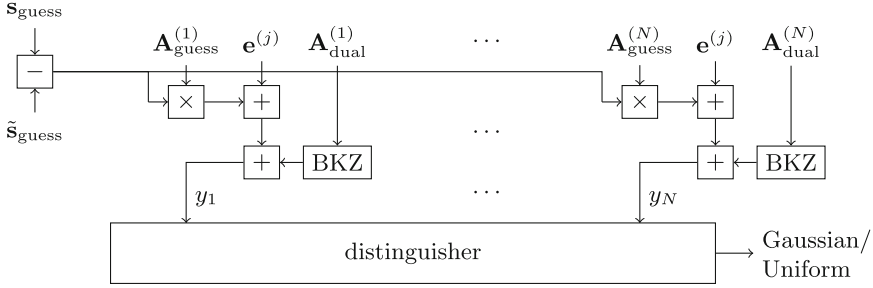


Fig. 1. Conceptual representation of the variables involved in Algorithm 1.

the distinguisher. However, computing a short vector is an expensive operation that we have to repeat many times. Another issue is that the attack requires an exponential number of LWE samples, something which is not always realistic.

As explained in the introduction, a series of work have progressively introduced the idea of generating *all short vectors* from a limited number of LWE sample, *i.e.* a single (\mathbf{A}, \mathbf{b}) . This is the case in [7, 23, 25], and [32] and it dramatically reduces the complexity of the attack. Unfortunately, the statistical analysis of these attacks has been lacking in the literature: [7, 23]⁴ and [25] offer no real proof of correctness to speak of. Only [32] tries to provide a complete proof of correctness, which is very detailed, but has to rely on statistical assumptions. Those assumptions have been called into question [20], and more importantly are extremely difficult to verify. Stepping back, we believe that the reason for this situation is that they try to analyse their attacks using a similar proof strategy to that of our basic dual attack (Sect. 3). However, the basic dual attack requires the independence of many variables to work. Since those variables become dependent in their attack, these papers inevitably have to assume or prove that non-independent quantities are “independent enough”.

In this section, we start completely from scratch: we design and analyze without any assumption a modern dual attack. Our proof scheme is completely different from the basic one and shows that those attacks do work. The main outcome of this proof is that we can finally understand the constraints on the various parameters that are necessary for the attack to work.

4.1 Intuition

Fix $\mathbf{s} \in \mathbb{Z}_q^n$ an unknown secret and (\mathbf{A}, \mathbf{b}) some LWE samples. Recall that $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ for some unknown $\mathbf{e} \in \mathbb{Z}_q^m$. As in the basic dual attack, we split the secret \mathbf{s} into two parts $\mathbf{s}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ and $\mathbf{s}_{\text{dual}} \in \mathbb{Z}_q^{n_{\text{dual}}}$ where $n = n_{\text{guess}} + n_{\text{dual}}$. The matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is correspondingly split into two parts:

$$\mathbf{A} = [\mathbf{A}_{\text{guess}} \ \mathbf{A}_{\text{dual}}], \quad \mathbf{s} = \begin{bmatrix} \mathbf{s}_{\text{guess}} \\ \mathbf{s}_{\text{dual}} \end{bmatrix}. \tag{3}$$

⁴ Part of [23] formally analyzes a similar attack to our basic attack. This paragraph only applies to the rest that relies on sieving to produce many short vectors.

The algorithm now makes a guess $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ on the value of $\mathbf{s}_{\text{guess}}$ and tries to check whether this guess is correct. Check that

$$\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}} = \mathbf{A}_{\text{guess}} \cdot (\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}}) + \mathbf{A}_{\text{dual}} \cdot \mathbf{s}_{\text{dual}} + \mathbf{e}. \quad (4)$$

Consider the lattice

$$L_q^\perp(\mathbf{A}_{\text{dual}}) = \{ \mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \mathbf{A}_{\text{dual}} = \mathbf{0} \bmod q \}. \quad (5)$$

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

$$g_W(\mathbf{x}) = \frac{1}{N} \sum_{j=1}^N \cos(2\pi \langle \mathbf{x}, \mathbf{w}_j \rangle / q) \quad (6)$$

for all $\mathbf{x} \in \mathbb{R}^m$. We will evaluate g_W at $\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}}$ for all $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ and keep the highest value. We now explain the intuition for this. Let $L = L_q(\mathbf{A}_{\text{dual}})$ to simplify notations. Recall that in Sect. 2.2, we have defined the standard periodic Gaussian function $f_{L,1/s}(\mathbf{x}) = \frac{\rho_{1/s}(\mathbf{x}+L)}{\rho_{1/s}(L)}$ for any $\mathbf{x} \in \mathbb{R}^m$ and $s > 0$. The important fact is that for large N , with high probability on the choice of the \mathbf{w}_j , g_W and $f_{L,1/s}$ are close everywhere for integer vectors (Lemma 6). This fact essentially comes from [5]. Therefore, it suffices to analyse the behaviour of $f_{L,1/s}$. For this, we rely on standard Gaussian tailbounds (Lemma 7) to get that for any $s > 0$ and $\mathbf{x} \in \mathbb{R}^m$, we essentially have

$$f_{L,1/s}(\mathbf{x}) \approx \rho_{1/s}(\text{dist}(\mathbf{x}, L)). \quad (7)$$

In other words, $f_{L,1/s}$ measures the *distance to the lattice* L .

We are now ready to see what makes the attack work. The intuition is that for most choices of \mathbf{A} and \mathbf{e} , for all $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}} \setminus \{ \mathbf{s}_{\text{guess}} \}$,

$$\text{dist}(\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \mathbf{s}_{\text{guess}}, L) < \text{dist}(\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}}, L) \quad (8)$$

and therefore

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

and the same will be true for g_W , which means that the algorithm will correctly output $\mathbf{s}_{\text{guess}}$. This is the main idea of our analysis but making it formal requires some care. The first step (Lemma 8) is to show that essentially

$$\text{if } 2 \|\mathbf{e}\| < \lambda_1(L_q(\mathbf{A})) \text{ then } f_{L, \frac{1}{s}}(\mathbf{e}) > f_{L, \frac{1}{s}}(\mathbf{e} + \mathbf{x}) \text{ for all } \mathbf{x} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L. \quad (9)$$

This requires some explanations. Going back to (8), we have that

$$\begin{aligned} \text{dist}(\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \mathbf{s}_{\text{guess}}, L) &= \text{dist}(\mathbf{e} + \mathbf{A}_{\text{dual}} \cdot \mathbf{s}_{\text{dual}}, L) \\ &= \text{dist}(\mathbf{e}, L) && \text{since } \mathbf{A}_{\text{dual}} \cdot \mathbf{s}_{\text{dual}} \in L \\ &= \|\mathbf{e}\| && \text{if } \|\mathbf{e}\| < \lambda_1(L)/2. \end{aligned}$$

On the other hand, if $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$ then

$$\begin{aligned} \text{dist}(\mathbf{b} - \mathbf{A}_{\text{guess}} \cdot \tilde{\mathbf{s}}_{\text{guess}}, L) &= \text{dist}(\mathbf{e} + \mathbf{A}_{\text{dual}} \cdot \mathbf{s}_{\text{dual}} + \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}}), L) \\ &= \text{dist}(\mathbf{e} + \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}}), L) \quad \text{since } \mathbf{A}_{\text{dual}} \cdot \mathbf{s}_{\text{dual}} \in L \\ &= \text{dist}(\mathbf{e} + \mathbf{x}, L) \end{aligned}$$

where

$$\mathbf{x} = \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}}) \in L_q(\mathbf{A}_{\text{guess}}).$$

Assume for now that $\mathbf{x} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L$ which we will see below is not always true but holds with probability exponentially close to 1 over the choice of \mathbf{A} . Then

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

The last step holds because $\mathbf{y} + \mathbf{z} \neq \mathbf{0}$ for all $\mathbf{z} \in L$ and $\mathbf{y} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L$. This is where our assumption that $\mathbf{x} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L$ is crucial. The condition in (8) now becomes

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

and this gives us (9) because $L + L_q(\mathbf{A}_{\text{guess}}) = L_q(\mathbf{A}_{\text{dual}}) + L_q(\mathbf{A}_{\text{guess}}) = L_q(\mathbf{A})$.

Now that we have (9), the second step is to apply it to \mathbf{A} . Recall that we made a crucial assumption above: it only applies to $\mathbf{e} + \mathbf{x}$ for $\mathbf{x} \in L_q(\mathbf{A}_{\text{guess}}) \setminus L$ where $\mathbf{x} = \mathbf{A}_{\text{guess}}(\mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}})$ and $\mathbf{s}_{\text{guess}} \neq \tilde{\mathbf{s}}_{\text{guess}}$. This condition is equivalent to $\mathbf{x} \notin \mathbf{A}_{\text{dual}}\mathbb{Z}_q^{n_{\text{dual}}} + q\mathbb{Z}^m$ since $L = L_q(\mathbf{A}_{\text{dual}})$. A sufficient condition for this to hold is that \mathbf{A} has full rank over \mathbb{Z}_q which happens with probability exponentially close to 1 over the choice of \mathbf{A} . This allows us to conclude (Theorem 5) that Algorithm 2, which essentially performs the steps highlighted above, works for almost all \mathbf{A} and \mathbf{e} that satisfy roughly $2\|\mathbf{e}\| < \lambda_1(L_q(\mathbf{A}))$. At this point, one can make two interesting observations:

- It tells us that if $2\|\mathbf{e}\| < \lambda_1(L_q(\mathbf{A}))$ then we can distinguish \mathbf{e} from any $\mathbf{e} + \mathbf{x}$ by using $f_{L,1/s}$. This makes intuitive sense since this condition guarantees that \mathbf{e} is the closest vector to $\mathbf{0}$ in $L_q(\mathbf{A})$ which is a *necessary condition* for the algorithm to work unconditionally⁵
- Even though we take short vectors in the dual lattice $L_q(\mathbf{A}_{\text{dual}})$, it looks like only the length of the shortest vectors in \mathbf{A} matters for the analysis! This is just a result of the simplifications that we have made above to give the intuition. The length of the dual vectors does play a role in Lemma 8 and the subsequent lemmas.

⁵ This condition could be relaxed if we allow the algorithm to fail for a small fraction of \mathbf{e} but this is out of the scope of this article.

4.2 Formal Analysis

This section gives a formal analysis of the intuitions from the previous section. We will reuse the notation defined there. Our first lemma formalizes that g_W , defined in (6) and used in the algorithm to compute the “score” of a guess, is very close to the periodic Gaussian function $f_{L_q(\mathbf{A}_{\text{dual}})}$.

Lemma 6. *Let $\mathbf{B} \in \mathbb{Z}_q^{m \times n}$, $s, \delta > 0$ and $N \in \mathbb{N}$. With probability at least $1 - q^m \cdot 2^{-\Omega(N\delta^2)}$ over the choice of $W = (\mathbf{w}_1, \dots, \mathbf{w}_N)$ from $D_{L_q^\perp(\mathbf{B}), qs}^N$, we have $|g_W(\mathbf{x}) - f_{L_q(\mathbf{B}), 1/s}(\mathbf{x})| \leq \delta$ for all $\mathbf{x} \in \mathbb{Z}^m$, where g_W is defined in (6) and $f_{L_q(\mathbf{B})}$ is defined in Sect. 2.2.*

Proof. Let $L = L_q(\mathbf{B})$ and for any j , let $\mathbf{w}'_j = \frac{1}{q}\mathbf{w}_j$ and $W' = (\mathbf{w}'_j)_j$. Since $\widehat{L} = \frac{1}{q}L_q^\perp(\mathbf{B})$ and $D_{L_q(\mathbf{B}), qs} = D_{q\widehat{L}, qs} = D_{\widehat{L}, s}$, we indeed have that W' is sampled from $D_{\widehat{L}, s}^N$ which is a probability distribution over \widehat{L} . Let $h = f_{L, 1/s}$ which is L -periodic, then $\widehat{h} = D_{\widehat{L}, s}$ by Lemma 3. For any $\mathbf{x} \in \mathbb{R}^m$, $g_W(\mathbf{x}) = h_{W'}(\mathbf{x})$ where $h_{W'}$ is defined in Lemma 5. Apply Lemma 5 to h with $X = \{0, \dots, q-1\}^m$ to get that with probability at least $1 - |X|2^{-\Omega(N\delta^2)}$ over the choice of W' , we have $|h(\mathbf{x}) - h_{W'}(\mathbf{x})| \leq \delta$ for all $\mathbf{x} \in L + X$. But $L = L_q(\mathbf{B})$ is a q -ary lattice, i.e. $q\mathbb{Z}^m \subset L$ so $L + X \supset q\mathbb{Z}^m + \{0, \dots, q-1\}^m = \mathbb{Z}^m$ which concludes. \square

The next lemma formalizes the idea that the periodic Gaussian function f_L estimates the distance of its argument and the lattice L .

Lemma 7. *Let $L \subset \mathbb{R}^m$ and $s > 0$, then for any $\mathbf{x} \in \mathbb{R}^m$:*

- $f_{L, 1/s}(\mathbf{x}) \geq \rho_{1/s}(\text{dist}(\mathbf{x}, L))$,
- if $\text{dist}(\mathbf{x}, L) \geq \tau := \frac{1}{s}\sqrt{m/2\pi}$ then $f_{L, 1/s}(\mathbf{x}) \leq \rho_{1/s}(\text{dist}(\mathbf{x}, L) - \tau)$.

Proof. The first fact is a direct consequence of Lemma 1. Indeed, write $\mathbf{x} = \mathbf{z} + \mathbf{t}$ where $\mathbf{z} \in L$ and $\mathbf{t} \in \mathbb{R}^m$ are such that $\text{dist}(\mathbf{x}, L) = \|\mathbf{t}\|$. Since $f_{L, 1/s}$ is L -periodic and $\mathbf{z} \in L$, $f_{L, 1/s}(\mathbf{x}) = f_{L, 1/s}(\mathbf{x} - \mathbf{z}) = f_{L, 1/s}(\mathbf{t}) \geq \rho_{1/s}(\mathbf{t}) = \rho_{1/s}(\|\mathbf{t}\|)$. For the second fact, let $\ell = \text{dist}(\mathbf{x}, L)$ and observe that by definition $(L - \mathbf{x}) \setminus B_m(\ell) = L - \mathbf{x}$. By assumption, $\ell \geq \tau := \frac{1}{s}\sqrt{m/2\pi}$, so we can apply Corollary 1 to get that $\rho_{1/s}((L - \mathbf{x}) \setminus B_m(\ell)) \leq \rho_{1/s}(\ell - \tau)\rho_{1/s}(L)$ and therefore

$$f_{L, 1/s}(\mathbf{x}) = \frac{\rho_{1/s}(L - \mathbf{x})}{\rho_{1/s}(L)} = \frac{\rho_{1/s}((L - \mathbf{x}) \setminus B_m(\ell))}{\rho_{1/s}(L)} \leq \rho_{1/s}(\ell - \tau).$$

\square

Lemma 8. *Let $\mathbf{B} \in \mathbb{Z}_q^{m \times n}$, $L \subset \mathbb{Z}^m$ a lattice, $\mathbf{e} \in \mathbb{Z}^m$, $s, \delta > 0$ and $N \in \mathbb{N}$. Let $\tau = \frac{1}{s}\sqrt{m/2\pi}$ and $\eta \geq 0$ and assume that $\lambda_1(L + L_q(\mathbf{B})) \geq \tau + \|\mathbf{e}\|$ and*

$$\rho_{1/s}(\mathbf{e}) - \rho_{1/s}(\lambda_1(L + L_q(\mathbf{B})) - \|\mathbf{e}\| - \tau) > 2\delta + \eta.$$

Then, with probability at least $1 - q^m \cdot 2^{-\Omega(N\delta^2)}$ over the choice of $W = (\mathbf{w}_1, \dots, \mathbf{w}_N)$ from $D_{L_q^\perp(\mathbf{B}), qs}^N$, we have

$$g_W(\mathbf{e}) \geq \rho_{1/s}(\mathbf{e}) - \delta > \rho_{1/s}(\lambda_1(L + L_q(\mathbf{B})) - \|\mathbf{e}\| - \tau) + \delta + \eta \geq g_W(\mathbf{e} + \mathbf{x}) + \eta$$

for all $\mathbf{x} \in L \setminus L_q(\mathbf{B})$, where g_W is defined in (6).

Proof. Apply Lemma 6 to get that with probability at least $1 - q^m \cdot 2^{-\Omega(N\delta^2)}$ over the choice of $\mathbf{w}_1, \dots, \mathbf{w}_N$ i.i.d. from $D_{L_q^\perp(\mathbf{B}), qs}$, we have $|g_W(\mathbf{y}) - f_{L_q(\mathbf{B}), 1/s}(\mathbf{y})| \leq \delta$ for all $\mathbf{y} \in \mathbb{Z}^m$. By Lemma 7, we have $g_W(\mathbf{e}) \geq f_{L_q(\mathbf{B}), 1/s}(\mathbf{e}) - \delta \geq \rho_{1/s}(\mathbf{e}) - \delta$.

Let $\mathbf{x} \in L \setminus L_q(\mathbf{B})$, then $\mathbf{z} - \mathbf{x} \in L + L_q(\mathbf{B})$ and $\mathbf{z} - \mathbf{x} \neq \mathbf{0}$ for any $\mathbf{z} \in L_q(\mathbf{B})$. As a result, $L_q(\mathbf{B}) - \mathbf{x} \subseteq (L + L_q(\mathbf{B})) \setminus \{\mathbf{0}\}$. Hence,

$$\text{dist}(\mathbf{x}, L_q(\mathbf{B})) = \min_{\mathbf{z} \in L_q(\mathbf{B})} \|\mathbf{x} + \mathbf{z}\| \geq \min_{\mathbf{y} \in (L + L_q(\mathbf{B})) \setminus \{\mathbf{0}\}} \|\mathbf{y}\| = \lambda_1(L + L_q(\mathbf{B})) \geq \tau + \|\mathbf{e}\|. \quad (10)$$

But then

$$\text{dist}(\mathbf{e} + \mathbf{x}, L_q(\mathbf{B})) \geq \text{dist}(\mathbf{x}, L_q(\mathbf{B})) - \|\mathbf{e}\| \geq \tau. \quad (11)$$

We can therefore apply Lemma 7 to get that for any $\mathbf{x} \in L \setminus \{\mathbf{0}\}$,

$$g_W(\mathbf{e} + \mathbf{x}) \leq f_{L_q(\mathbf{B}), 1/s}(\mathbf{e} + \mathbf{x}) + \delta \leq \rho_{1/s}(\text{dist}(\mathbf{e} + \mathbf{x}, L_q(\mathbf{B})) - \tau) + \delta.$$

Since $\rho_{1/s} : [0, \infty) \rightarrow \mathbb{R}$ is decreasing, and reusing (10) and (11) we further have

$$\begin{aligned} \rho_{1/s}(\text{dist}(\mathbf{e} + \mathbf{x}, L_q(\mathbf{B})) - \tau) &\leq \rho_{1/s}(\text{dist}(\mathbf{x}, L_q(\mathbf{B})) - \|\mathbf{e}\| - \tau) \\ &\leq \rho_{1/s}(\lambda_1(L + L_q(\mathbf{B})) - \|\mathbf{e}\| - \tau). \end{aligned}$$

Putting everything together, and using our assumption, we have

$$g_W(\mathbf{e}) - g_W(\mathbf{e} + \mathbf{x}) \geq \rho_{1/s}(\mathbf{e}) - \rho_{1/s}(\lambda_1(L + L_q(\mathbf{B})) - \|\mathbf{e}\| - \tau) - 2\delta > \eta$$

□

We can now state our main result by putting everything together. It will be useful to note that $L_q(\mathbf{A}_{\text{guess}}) + L_q(\mathbf{A}_{\text{dual}}) = L_q(\mathbf{A})$ which is readily verified.

Algorithm 2: Modern dual attack

Input: $m, n = n_{\text{guess}} + n_{\text{dual}}$ (see (1)), q prime power, $N \in \mathbb{N}$

Input: LWE sample (\mathbf{A}, \mathbf{b}) , list $W = (\mathbf{w}_1, \dots, \mathbf{w}_N)$ of vectors in $L_q^\perp(\mathbf{A}_{\text{dual}})$.

Output: (Guess of) the first n_{guess} coordinates of the secret, or \perp .

1 $\mathbf{s}_{\text{guess}} \leftarrow \perp; S_{\text{max}} \leftarrow 0$;

2 **for** $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ **do**

3 Compute the list y_1, \dots, y_N where $y_j = \mathbf{w}_j^T(\mathbf{b} - \mathbf{A}_{\text{guess}}\tilde{\mathbf{s}}_{\text{guess}})$;

4 $S \leftarrow \sum_{j=1}^N \cos(2\pi y_j/q)$;

5 **if** $S \geq S_{\text{max}}$ **then** $S_{\text{max}} \leftarrow S; \mathbf{s}_{\text{guess}} \leftarrow \tilde{\mathbf{s}}_{\text{guess}}$;

6 **return** $\mathbf{s}_{\text{guess}}$

Theorem 5. Let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, $\mathbf{e} \in \mathbb{Z}^m$, $\mathbf{s} \in \mathbb{Z}_q^n$, $s, \delta > 0$ and $N \in \mathbb{N}$. Let $\tau = \frac{1}{s} \sqrt{m/2\pi}$. Assume that $m \geq n$, \mathbf{A} has full rank, $\lambda_1(L_q(\mathbf{A})) \geq \tau + \|\mathbf{e}\|$, and

$$\rho_{1/s}(\mathbf{e}) - \rho_{1/s}(\lambda_1(L_q(\mathbf{A}))) - \|\mathbf{e}\| - \tau > 2\delta.$$

Let $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$. Let $W = (\mathbf{w}_1, \dots, \mathbf{w}_N)$ be samples from $D_{L_q^\perp(\mathbf{A}_{\text{dual}}), qs}^N$, then Algorithm 2 on $(m, n_{\text{guess}}, n_{\text{dual}}, q, N, (\mathbf{A}, \mathbf{b}), W)$ runs in time $\text{poly}(m, n) \cdot (N + q^{n_{\text{guess}}})$ and returns $\mathbf{s}_{\text{guess}}$ with probability at least $1 - q^m \cdot 2^{-\Omega(N\delta^2)}$ over the choice of W .

Proof. Let $\mathbf{B} = \mathbf{A}_{\text{dual}}$ and $L = L_q(\mathbf{A}_{\text{guess}})$. Then $L + L_q(\mathbf{B}) = L_q(\mathbf{A})$. Our assumptions are therefore exactly that of Lemma 8 for $\eta = 0$ which we can apply to get that with probability at least $1 - q^m \cdot 2^{-\Omega(N\delta^2)}$ over the choice of $W = (\mathbf{w}_1, \dots, \mathbf{w}_N)$ from $D_{L_q^\perp(\mathbf{B}), qs}^N = D_{L_q^\perp(\mathbf{A}_{\text{dual}}), qs}^N$, we have

$$g_W(\mathbf{e}) > g_W(\mathbf{e} + \mathbf{x}) \quad (12)$$

for all $\mathbf{x} \in L \setminus L_q(\mathbf{A}_{\text{dual}})$, where g_W is defined in (6). Furthermore, \mathbf{A} has full rank and $m \geq n$ so its columns are linearly independent over \mathbb{Z}_q and

$$L \setminus L_q(\mathbf{A}_{\text{dual}}) = L_q(\mathbf{A}_{\text{guess}}) \setminus L_q(\mathbf{A}_{\text{dual}}) = L_q(\mathbf{A}_{\text{guess}}) \setminus q\mathbb{Z}^m. \quad (13)$$

Assume that we are in the case where W satisfies the above inequalities and consider the run of Algorithm 2 on $(m, n_{\text{guess}}, n_{\text{dual}}, q, N, (\mathbf{A}, \mathbf{b}), W)$. The algorithm tests all possible values of $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ and returns the one that maximizes S . Let $\tilde{\mathbf{s}}_{\text{guess}} \in \mathbb{Z}_q^{n_{\text{guess}}}$ and $\Delta\tilde{\mathbf{s}}_{\text{guess}} = \mathbf{s}_{\text{guess}} - \tilde{\mathbf{s}}_{\text{guess}}$. First note that

$$\begin{aligned} \mathbf{b} - \mathbf{A}_{\text{guess}}\tilde{\mathbf{s}}_{\text{guess}} &= (\mathbf{A}\mathbf{s} + \mathbf{e} \bmod q) - \mathbf{A}_{\text{guess}}\tilde{\mathbf{s}}_{\text{guess}} \\ &= \mathbf{A}_{\text{dual}}\mathbf{s}_{\text{dual}} + \mathbf{A}_{\text{guess}}\Delta\tilde{\mathbf{s}}_{\text{guess}} + \mathbf{e} \bmod q. \end{aligned}$$

For any j , let $y_j(\tilde{\mathbf{s}}_{\text{guess}})$ be the value computed at Line 3. Note that

$$\begin{aligned} y_j(\tilde{\mathbf{s}}_{\text{guess}}) &= \mathbf{w}_j^T(\mathbf{b} - \mathbf{A}_{\text{guess}}\tilde{\mathbf{s}}_{\text{guess}}) \\ &= \mathbf{w}_j^T \mathbf{A}_{\text{dual}}\mathbf{s}_{\text{dual}} + \mathbf{w}_j^T(\mathbf{A}_{\text{guess}}\Delta\tilde{\mathbf{s}}_{\text{guess}} + \mathbf{e}) \bmod q \end{aligned}$$

but $\mathbf{w}_j \in L_q^\perp(\mathbf{A}_{\text{dual}})$ so $\mathbf{w}_j^T \mathbf{A}_{\text{dual}} = \mathbf{0} \bmod q$, hence

$$= \mathbf{w}_j^T(\mathbf{A}_{\text{guess}}\Delta\tilde{\mathbf{s}}_{\text{guess}} + \mathbf{e}) \bmod q.$$

Let $S(\tilde{\mathbf{s}}_{\text{guess}})$ be the value computed at Line 4 and check that

$$\begin{aligned} S(\tilde{\mathbf{s}}_{\text{guess}}) &= \sum_{j=1}^N \cos(2\pi y_j(\tilde{\mathbf{s}}_{\text{guess}})/q) \\ &= \sum_{j=1}^N \cos(2\pi \mathbf{w}_j^T(\mathbf{A}_{\text{guess}}\Delta\tilde{\mathbf{s}}_{\text{guess}} + \mathbf{e})/q) \quad \text{by periodicity of cos} \\ &= Ng_W(\mathbf{A}_{\text{guess}}\Delta\tilde{\mathbf{s}}_{\text{guess}} + \mathbf{e}). \end{aligned}$$

There are two cases to distinguish:

- If $\tilde{\mathbf{s}}_{\text{guess}} = \mathbf{s}_{\text{guess}}$ then $S(\tilde{\mathbf{s}}_{\text{guess}}) = Ng_W(\mathbf{e})$.
- If $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$ then $S(\tilde{\mathbf{s}}_{\text{guess}}) = Ng_W(\mathbf{e} + \mathbf{x})$ where $\mathbf{x} = \mathbf{A}_{\text{guess}}\Delta\tilde{\mathbf{s}}_{\text{guess}} \in L_q(\mathbf{A}_{\text{guess}}) = L$. But \mathbf{A} (and hence $\mathbf{A}_{\text{guess}}$) has full rank by assumption and $\Delta\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{0}$ so $\mathbf{x} \neq \mathbf{0} \pmod q$. It follows by (13), $\mathbf{x} \in L_q(\mathbf{A}_{\text{dual}}) \setminus q\mathbb{Z}^m = L \setminus L_q(\mathbf{A}_{\text{dual}})$. Hence, by (12), $S(\tilde{\mathbf{s}}_{\text{guess}}) < Ng_W(\mathbf{e}) = S(\mathbf{s}_{\text{guess}})$.

This shows that $S(\mathbf{s}_{\text{guess}}) > S(\tilde{\mathbf{s}}_{\text{guess}})$ for all $\tilde{\mathbf{s}}_{\text{guess}} \neq \mathbf{s}_{\text{guess}}$. Therefore, Algorithm 2 correctly returns $\mathbf{s}_{\text{guess}}$. Note that the entire argument was under the assumption that (12) holds for W , which we already argued holds with probability at least $1 - q^m \cdot 2^{-\Omega(N\delta^2)}$.

The naive analysis of the complexity is straightforward and gives $q^{n_{\text{guess}}} \cdot \text{poly}(m, n) \cdot N$. By using the DFT trick as we did in the proof of Theorem 4 we can improve the running time to $\text{poly}(m, n) \cdot (N + q^{n_{\text{guess}}})$.

4.3 Informal Application

Choosing the parameters in order to apply Theorem 5 is not immediately obvious. In this section, we explain how to do so in a concrete case of interest. In order to simplify things, we will neglect some factors and point out the various lemmas that can be used to make this reasoning completely formal.

Fix n, m and let q be a prime power. Let $\mathbf{s} \in \mathbb{Z}_q^n$ be a secret and $\sigma_e > 0$. Let (\mathbf{A}, \mathbf{b}) be sampled from $\text{LWE}(m, \mathbf{s}, D_{\mathbb{Z}_q, \sigma_e})$, and \mathbf{e} so that $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$. By Corollary 1, we have

$$\|\mathbf{e}\| \lesssim \sigma_e \sqrt{m/2\pi} \tag{14}$$

with high probability. Let $s > 0$ to be defined later. We choose δ to be quite smaller than the smallest possible value $\rho_{1/s}(\|\mathbf{e}\|)$, for example

$$\delta = \frac{1}{100} \rho_{1/s}(\sigma_e \sqrt{m/2\pi}) = \frac{1}{100} e^{-ms^2\sigma_e^2/2}. \tag{15}$$

We choose N accordingly so that the success probability is very high, *i.e.*

$$N = \frac{\text{poly}(m) + n \log_2(q)}{\delta^2}. \tag{16}$$

\mathbf{A} has full rank with high probability and therefore $\det(L_q(\mathbf{A})) = q^{m-n}$. By Theorem 3, and the informal Corollary 2, we have

$$\lambda_1(L_q(\mathbf{A})) \gtrsim GH(L_q(\mathbf{A})) = \text{vol}(B_m)^{-1/n} q^{1-m/n} \approx \sqrt{\frac{m}{2\pi e}} q^{1-n/m}.$$

Let $\tau = \frac{1}{s} \sqrt{m/2\pi}$. In order to apply Theorem 5, we need to satisfy the conditions

$$\lambda_1(L_q(\mathbf{A})) \geq \tau + \|\mathbf{e}\| \quad \text{and} \quad \rho_{1/s}(\mathbf{e}) - \rho_{1/s}(\lambda_1(L_q(\mathbf{A})) - \|\mathbf{e}\| - \tau) > 2\delta.$$

Since we have chosen δ to be very small compared to $\rho_{1/s}(\mathbf{e})$, those inequalities can be shown (see [39, Appendix D]) to be essentially equivalent to

$$\lambda_1(L_q(\mathbf{A})) \geq \tau + 2\|\mathbf{e}\|.$$

This condition will be satisfied when $\sqrt{\frac{m}{2\pi e}}q^{1-n/m} \geq \frac{1}{s}\sqrt{m/2\pi} + 2\sigma_e\sqrt{m/2\pi}$ that is

$$q^{1-n/m} \geq \left(\frac{1}{s} + 2\sigma_e\right)\sqrt{e}. \quad (17)$$

In other words, we have a lower bound on s . We observe that there is a trade-off between the cost of sampling from $D_{L_q^\perp(\mathbf{A}_{\text{dual}}),qs}$ and the cost of running Algorithm 2 since a large value of s :

- makes it easy to sample from $D_{L_q^\perp(\mathbf{A}_{\text{dual}}),qs}$,
- but makes $\delta = \frac{1}{100}\rho_{1/s}(\sigma_e\sqrt{m/2\pi})$ small and therefore $N = \Omega(\delta^{-2})$, and the complexity, gigantic.

We note that the total complexity of the attack, including the cost of generating the small dual vectors, is a highly nontrivial function of the parameters. Consequently, it is not at all clear that the optimal choice of s is the lower bound identified above. We will analyze the complexity in greater detail in the next section.

4.4 Complexity Estimates

In this section, we describe how to concretely estimate the complexity of the attack described in Sect. 4 and provide numbers for Kyber. We continue with the setup from the previous section (Sect. 4.3) which we do not repeat. Recall that by Theorem 5, the complexity of the attack, to which we add the cost $T_{\text{sampling}}(N, qs)$ of sampling N independent Gaussian vectors according to $D_{L_q^\perp(\mathbf{A}_{\text{dual}}),qs}$ is

$$\text{poly}(m, n) \cdot (N + q^{n_{\text{guess}}}) + T_{\text{sampling}}(N, qs) \quad (18)$$

and it succeeds with very high probability given the choice of the parameters above. For the sampling of the dual vectors, we propose the following approach: given a block size $2 \leq \beta \leq m$,

1. compute a basis of $L_q^\perp(\mathbf{A}_{\text{dual}})$,
2. run BKZ with block size β on this basis to obtain a reduced basis \mathbf{B} ,
3. use the Markov chain Monte Carlo (MCMC) based Gaussian sampler from [46] (Theorem 1) for parameter qs with basis \mathbf{B} to generate N independent samples.

The complexity of this procedure is

$$T_{\text{sampling}}(N) = T_{\text{BKZ}}(m, \beta) + N \cdot T_{\text{MCMC}}(L_q^\perp(\mathbf{A}_{\text{dual}}), qs) \quad (19)$$

where $T_{\text{BKZ}}(m, \beta)$ is the cost of BKZ and $T_{\text{MCMC}}(L, s)$ is the cost of producing one sample from $D_{L,s}$. We apply Theorem 1 to get that

$$T_{\text{MCMC}}(L_q^\perp(\mathbf{A}_{\text{dual}}), qs) = \ln\left(\frac{1}{\varepsilon}\right) \cdot \frac{1}{\Delta} \cdot \text{poly}(n), \quad \Delta = \frac{\rho_{qs}(L_q^\perp(\mathbf{A}_{\text{dual}}))}{\prod_{i=1}^n \rho_{qs}/\|\tilde{\mathbf{b}}_i\|(\mathbb{Z})} \quad (20)$$

where $\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_n$ are the Gram-Schmidt vectors of the BKZ- β -reduced basis \mathbf{B} of $L_q^\perp(\mathbf{A}_{\text{dual}})$ and $\varepsilon > 0$. Note that the output distribution of the algorithm is ε -close to the discrete Gaussian. Since we are going to use N samples, and by the data processing inequality, this translates into a failure probability of $N\varepsilon$ for the algorithm, so we need to choose ε to be quite small, e.g. $\varepsilon \ll 1/N$. Putting (18) and (19) together we get that the total complexity of the attack is

$$\text{poly}(m, n) \cdot (N + q^{n_{\text{guess}}}) + T_{\text{BKZ}}(m, \beta) + N \cdot T_{\text{MCMC}}(L_q^\perp(\mathbf{A}_{\text{dual}}), qs) \quad (21)$$

subject to the constraints (14), (15), (16), (17) which we summarize below:

$$\begin{aligned} \delta &= \frac{1}{100} e^{-ms^2\sigma_e^2/2}, & N &= \frac{\text{poly}(m) + n \log_2(q)}{\delta^2}, \\ q^{1-n/m} &\geq (\frac{1}{s} + 2\sigma_e)\sqrt{e}, & \|\mathbf{e}\| &\lesssim \sigma_e \sqrt{m/2\pi}, \\ \varepsilon &\ll 1/N. \end{aligned}$$

In practice, computing Δ with (20) is nontrivial. One can show that (see [39, Appendix E])

$$\frac{1}{\Delta} \leq \prod_{i=1}^m \rho_{\|\tilde{\mathbf{b}}_i\|/qs}(\mathbb{Z}) \quad (22)$$

which is easier to estimate but still requires to estimate the $\|\tilde{\mathbf{b}}_i\|$. For this, we can assume that the Geometric Series Assumption (GSA) [42] holds for BKZ- β reduced basis. The GSA is known to be reasonably accurate when $\beta \ll m$ and $\beta \gg 50$ which is the case in our experiments, but it does not correctly model what happens in the last $m - \beta$ coordinates [1]. For our purpose, we consider the GSA to be enough to obtain credible estimates on the complexity.

Independently of the GSA, however, formula (22) is expensive to compute due the product of m terms. Indeed, we will need to compute this quantity many times in our optimizer to find a good set of parameters (see below). For the estimates below, we use (22) and the GSA to compute the final complexity estimate but we use the approximate formula below in the parameter optimizer which is very cheap to compute:

$$\text{if } \|\mathbf{b}_1\| \leq 2qs \text{ then } \frac{1}{\Delta} \lesssim \exp\left(\log(1 + 2e^{-\pi\alpha}) + \frac{2}{\ln(H_\beta^4)} E_1(\pi\alpha)\right) \quad (23)$$

where $\alpha = (qs)^2 / \|\mathbf{b}_1\|^2$, E_1 is the generalized exponential integral and H_β is the Hermite factor for BKZ- β reduced basis. See [39, Appendix E] for more details.

In order to find a good set of parameters, we wrote an optimizer that tries all reasonable values of m , β and n_{guess} , and sets s to

$$s = \max\left(\frac{\sqrt{e}}{q^{1-n/m} - 2\sqrt{e}\sigma_e}, \frac{\|\mathbf{b}_1\|}{2q}\right)$$

so that we can use (23). We also limit the range of m to $[\frac{3}{2}n, 2n]$ so that the ratio m/n is not too close to 0 and 1.

In Table 1, we give the complexity estimates of our algorithm, computed by our optimizer. The first set of columns corresponds to the algorithm analyzed above, including the use of the GSA to estimate the complexity of the Gaussian sampling as described in [39, Appendix E]. To estimate the complexity of BKZ, we use the cost estimates in [8, 32] using [13] as the sieving oracle; specifically, we rely on the “lattice estimator” of [9].

Those cost are not competitive with the state of the art because our algorithm does not include modulus switching. Modulus switching is a critical component to reduce the complexity but its formal analysis is nontrivial and therefore we decided not to include it in this paper. In order to get an idea of what our algorithm extended with modulus switching would give, we include a second set of columns where we simply replace $q^{n_{\text{guess}}}$ by $2^{n_{\text{guess}}}$ in (21) which would correspond to switching the modulus to 2 in the guessing part. We emphasize that this is only a very rough estimate and not a formal analysis. The real complexity with modulus switching will most likely be higher than what we report. Furthermore, all our complexity estimates ignore the polynomial factors.

Table 1. Dual attack cost estimates and their parameters as described in Sect. 4.4. All costs are logarithms in base two. Note that the cost of attacks with modulus switching are estimates of what an algorithm with modulus switching could give if the algorithm of Sect. 4 was extended with modulus switching.

Scheme	No modulus switching						With modulus switching					
	attack	m	n_{guess}	n_{dual}	β	s	attack	m	n_{guess}	n_{dual}	β	s
Kyber512	185	1013	15	497	550	0.200	141	763	141	371	390	0.170
Kyber768	273	1469	23	745	870	0.260	202	1169	201	567	610	0.240
Kyber1024	376	2025	31	993	1230	0.270	279	1575	261	763	890	0.260

5 Quantum Dual Attack

In this section, we present a quantum version of Algorithm 2 and show that we can obtain a speed-up on the complexity. The technique is inspired by [10] which was never published and is a quantum variant of [32].

5.1 Algorithm and Analysis

We will need a quantum algorithm which estimates the mean value of $\cos(2\pi(\langle \mathbf{w}_i, \mathbf{b} \rangle)/q)$ where the \mathbf{w}_i are vectors accessible via a quantum oracle. This mean value can be used to compute the DFT sums in the algorithm much faster than with a classical computer. The idea is inspired by [2, Theorem 47] and can be seen as a special case of quantum speedup of Monte Carlo methods [37]. For more background on quantum algorithms, we refer the readers to [10, Sections 2.4 and 4].

Theorem 6 ([10, Theorem 5]). *Let N be a positive integer and W be a list of N vectors in \mathbb{Z}^n : $\mathbf{w}_0, \dots, \mathbf{w}_{N-1}$. Let $f_W(\mathbf{b}) = \frac{1}{N} \sum_{i=0}^{N-1} \cos(2\pi(\langle \mathbf{w}_i, \mathbf{b} \rangle)/q)$, where $\mathbf{b} \in \mathbb{Z}_q^n$. Let \mathcal{O}_W be defined by $\mathcal{O}_W : |j\rangle|0\rangle \mapsto |j\rangle|\mathbf{w}_j\rangle$. For any $\epsilon, \delta > 0$, there exists a quantum algorithm \mathcal{A} that given $\mathbf{b} \in \mathbb{Z}_q^n$ and oracle access to \mathcal{O}_W outputs $\mathcal{A}^{\mathcal{O}_W}(\mathbf{b})$ which satisfies $|\mathcal{A}^{\mathcal{O}_W}(\mathbf{b}) - f_W(\mathbf{b})| \leq \epsilon$ with probability $1 - \delta$. The algorithm makes $\mathcal{O}(\epsilon^{-1} \cdot \log \frac{1}{\delta})$ queries to \mathcal{O}_W , and requires $O(\log(\frac{1}{\epsilon}) + \text{poly}(\log(n)))$ qubits.*

We will have to search for a minimum element in a collection but the oracle that computes the value of each element is probabilistic and may return a wrong result with small probability. We say that a (probabilistic) real function f has bounded error if there exists $x \in \mathbb{R}$ such that $f()$ returns x with probability at least $9/10$. The problem of finding the minimum in a collection (without errors) has been studied in [21, Theorem 1]. On the other hand, the problem of searching for a marked element in a collection with bounded-error oracle has been studied in [26]. This idea can easily be used to adapt the algorithm of [21] to bounded-error oracles. Indeed, the algorithm in [21] simply performs a constant number of Grover searches by marking nodes that are bigger than the current value. Therefore it suffices to replace this Grover search by the algorithm of [26].

Theorem 7 ([26]+[21]). *Given n algorithms, quantum or classical, each computing some real value with bounded error probability, there is a quantum algorithm that makes an expected $O(\sqrt{n})$ queries and with probability at least $9/10$ returns the index of the minimum among the n values. This algorithm uses $\text{poly}(\log(n))$ qubits.*

Algorithm 3: Quantum modern dual attack

- Input:** $m, n = n_{\text{guess}} + n_{\text{dual}}$ (see (1)), q prime power, $N \in \mathbb{N}, \eta > 0$.
 - Input:** LWE sample (\mathbf{A}, \mathbf{b}) .
 - Input:** Oracle \mathcal{O}_W for a list $W = (\mathbf{w}_1, \dots, \mathbf{w}_N)$ of vectors in $L_q^\perp(\mathbf{A}_{\text{dual}})$.
 - Output:** (Guess of) the first n_{guess} coordinates of the secret, or \perp .
 - 1 Use Theorem 6 to create an algorithm \mathcal{A} with $\delta = \frac{1}{10}, \epsilon = \eta$ and q ;
 - 2 **create oracle** $\hat{\mathcal{O}}(\tilde{\mathbf{s}}_{\text{guess}})$:
 - 3 | **return** $\mathcal{A}^{\mathcal{O}_W}(\mathbf{b} - \mathbf{A}_{\text{guess}}\tilde{\mathbf{s}}_{\text{guess}})$
 - 4 Use Theorem 7 to find $\tilde{\mathbf{s}}_{\text{guess}}$ such that $\hat{\mathcal{O}}(\tilde{\mathbf{s}}_{\text{guess}})$ is maximum ;
 - 5 **return** $\tilde{\mathbf{s}}_{\text{guess}}$
-

Theorem 8 ([39, Appendix F.1]). *Let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}, \mathbf{e} \in \mathbb{Z}^m, \mathbf{s} \in \mathbb{Z}_q^n, s, \delta > 0$ and $N \in \mathbb{N}$. Let $\tau = \frac{1}{s} \sqrt{m/2\pi}$ and $\eta > 0$. Assume that $m \geq n, \mathbf{A}$ has full rank, $\lambda_1(L_q(\mathbf{A})) \geq \tau + \|\mathbf{e}\|$, and*

$$\rho_{1/s}(\mathbf{e}) - \rho_{1/s}(\lambda_1(L_q(\mathbf{A})) - \|\mathbf{e}\| - \tau) > 2\delta + \eta.$$

Let $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q$. Let $W = (\mathbf{w}_1, \dots, \mathbf{w}_N)$ be samples from $D_{L_q^\perp(\mathbf{A}_{\text{dual}}), qs}^N$ and \mathcal{O}_W an oracle for W in the sense of Theorem 6. Then Algorithm 3 on $m, n_{\text{guess}}, n_{\text{dual}}, q, N, \eta/2, (\mathbf{A}, \mathbf{b}), \mathcal{O}_W$ makes an expected $O(\eta^{-1} \cdot q^{n_{\text{guess}}/2})$ calls to \mathcal{O}_W and returns $\mathbf{s}_{\text{guess}}$ with probability at least $1 - q^m \cdot 2^{-\Omega(N\delta^2)}$ over the choice of W . The algorithm uses $O(\log(\eta^{-1}) + \text{poly}(\log(N)))$ qubits.

In terms of proofs, the correctness of the quantum algorithm is very similar to the classical one. The main difference is that we use Theorem 6 to compute g_W which only returns an approximation. This adds an additional error term that we can take into account in Lemma 8 using η .

5.2 Applications

In order to apply Theorem 8, one needs to provide an oracle \mathcal{O}_W to access the samples. The implementation of this oracle has a significant impact on the complexity since it is queried an exponential number of times by the algorithm. We outline two possible implementations. Before that, note that in practice we will usually choose η to be a small value compared to δ in Theorem 8, say $\eta = \delta/100$. This way, η has almost no influence on the maximum length of the errors \mathbf{e} that we can handle.

BKZ Preprocessing with a Quantum Klein Sampler. For a value of s that is not too small, one can first compute a basis $L_q^\perp(\mathbf{A}_{\text{dual}})$ of the dual lattice and reduce it using BKZ with block size β to obtain a new basis \mathbf{M} . One then creates a quantum circuit that implements the Klein sampler [24] with \mathbf{M} hard-coded in the circuit. This circuit will be the oracle \mathcal{O}_W . In the details, the Klein sampler is a probabilistic algorithm so we can view it as a deterministic algorithm that takes random coins (and \mathbf{M}) as input. We can see the input j of the oracle as the value of the random coins so that the outputs $\mathbf{w}_1, \dots, \mathbf{w}_N$ that correspond to inputs $1, \dots, N$ are distributed according to the Gaussian distribution. Since the Klein sampler runs in polynomial time, each call to \mathcal{O}_W takes polynomial time. The BKZ preprocessing is purely classical and done only once before the quantum algorithm runs. This means that the total runtime will be⁶, per Theorem 8

$$T_{\text{BKZ}}(\beta) + \sqrt{N} \cdot q^{n_{\text{guess}}/2} \cdot \text{poly}(\log(m)).$$

This is always better than the classical complexity since $\sqrt{N \cdot q^{n_{\text{guess}}}} \leq N + q^{n_{\text{guess}}}$. Note that when using a Klein sampler, the value of s is a function of the quality of the basis \mathbf{M} and therefore depends on β . Furthermore, it is impossible for s to be smaller than the smoothing parameter of the lattice this way. Alternatively, one could also use the MCMC sampler that we used in Sect. 4.4: although its running time is not polynomial, it only uses polynomial memory so it would still only require a polynomial number of qubits, and it allows one to

⁶ We chose $\eta = \delta/100$ and we explained in Sect. 4.3 that we need to choose $N = 1/\delta^2$ up to polynomial factors so $\eta^{-1} = \text{poly}(\log(m)) \cdot \sqrt{N}$.

choose smaller values of s which seems to be quite beneficial. Note that in both cases (Klein and MCMC), we get no quantum speed up on the sampling.

Classical Sampler with a Quantum Memory. A feature of the Klein sampler is that it can output an arbitrary number of samples and the running time is proportional to the number of samples. This is not the case of all samplers. For example, [4] describes Gaussian samplers that works for smaller values of s than the smoothing parameter and produces $2^{n/2}$ samples but runs in time 2^n , even if we only require one sample. [3] contains another such algorithm with a time-space trade-off. Using such samplers with our quantum algorithm is problematic because the samples are produced and stored in a classical memory, but the algorithm requires quantum oracle access to those samples. We have two options:

- We can assume that we have access to a QRACM (classical memory with quantum random access) [29]. A QRACM of size N is a special quantum memory holding N classical values but providing $O(\log(N))$ -time quantum access to those values. Such a QRACM directly implements the oracle \mathcal{O}_W so the total execution time becomes

$$T_{\text{sampler}} + \sqrt{N} \cdot q^{n_{\text{guess}}/2} \cdot \log(N) \cdot \text{poly}(\log(m)).$$

We note however that practical realizability of QRACM is debated and is potentially a strong assumption. We refer the readers to [27] for more details.

- We can replace \mathcal{A} in the algorithm by a very large circuit containing all N hard-coded samples that computes the sum g_W in a naive way (without Theorem 6). This circuit will take time $N \text{poly}(\log(m))$ to evaluate, therefore the total complexity will be

$$T_{\text{sampler}} + N \cdot q^{n_{\text{guess}}/2} \cdot \text{poly}(\log(m)).$$

Note that this might be worse than the classical algorithm if the value of N is larger than $q^{n_{\text{guess}}/2}$.

Finally, we note that presently samplers such as [3] are still too expensive to be useful in dual attacks but future samplers might get more efficient.

6 Comparison with [20]’s Contradictory Regime

In [20], the authors claim that [32] falls into what they call the “contradictory regime” and conclude that the result is most likely incorrect. They similarly conclude the recent derivative works [10, 16], as well as [25] are flawed. They do so by reconstructing the key heuristic claim of [32] and showing, both by theoretical arguments and experiments, that this heuristic is incorrect. We copy this heuristic below, slightly adjusted to our notations. In the heuristic, the function f_W is the same as h_W in Lemma 5, which is the same as g_W defined in (6) up to a factor $1/q$ in the cosine.

Heuristic 1 ([20, Heuristic Claim 3]). Let $\Lambda \subseteq \mathbb{R}^n$ be a random lattice of determinant 1, $\mathcal{W} \subseteq \widehat{\Lambda}$ be the set consisting of the $N = (4/3)^{n/2}$ shortest vectors of $\widehat{\Lambda}$. For some $\sigma > 0$ and $T \geq 1$, consider $\mathbf{t}_{BDD} \leftarrow_s \mathcal{N}(0, \sigma^2)^n$ and i.i.d $\mathbf{t}_{unif}^{(i)} \leftarrow_s \mathcal{U}(\mathbb{R}^n/\Lambda)$ where $i \in \{1, \dots, T\}$. Let⁷ $\ell = \sqrt{4/3} \cdot \text{GH}(n)$, $\varepsilon = \exp(-2\pi^2\sigma^2\ell^2)$. If $\ln T \leq N\varepsilon^2$,

$$\Pr \left[f_{\mathcal{W}}(\mathbf{t}_{BDD}) > f_{\mathcal{W}}(\mathbf{t}_{unif}^{(i)}) \text{ for all } i \in \{1, \dots, T\} \right] \geq 1 - O\left(\frac{1}{\sqrt{\ln T}}\right)$$

where $\mathcal{N}(0, \sigma^2)$ denotes the normal distribution.

There are several obvious (minor) problems about this heuristic since [32] works with integer lattices and discrete Gaussians. As a first step, we rewrite this heuristic in a way that is closer to [32] and we also change the notations to ours (see [39, Appendix G.1] for details about the rewrite).

Heuristic 2 ([20, Heuristic Claim 3] adapted). Let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ with i.i.d. coefficients. Let $L = L_q(\mathbf{A}) \subseteq \mathbb{Z}^m$ and $W \subseteq L_q^\perp(\mathbf{A})$ be the set consisting of the $N = (4/3)^{d/2}$ shortest vectors of $L_q^\perp(\mathbf{A})$. For some $\sigma_e > 0$ and $T \geq 1$, consider $\mathbf{e} \leftarrow_s D_{\mathbb{Z}_q, \sigma_e}^n$ and i.i.d $\mathbf{t}_{unif}^{(i)} \leftarrow_s \mathcal{U}(\mathbb{Z}^m/L)$ where $i \in \{1, \dots, T\}$. Let $\ell = \sqrt{4/3} \cdot \text{GH}(L)$, $\varepsilon = \exp(-\pi\sigma_e^2\ell^2)$. If $\ln T \leq N\varepsilon^2$,

$$\Pr \left[g_W(\mathbf{e}) > g_W(\mathbf{t}_{unif}^{(i)}) \text{ for all } i \in \{1, \dots, T\} \right] \geq 1 - O\left(\frac{1}{\sqrt{\ln T}}\right).$$

In [20, Section 4.2 and 4.3], the authors argue by theoretical arguments that Heuristic 1 does not hold. Although [20] did not define what they mean by “random lattice” in the heuristic, they in fact use random q -ary lattices in their experiments and also the theoretical properties of “random lattices” that they use hold for q -ary lattices. Therefore, their analysis holds also for Heuristic 2.

Their reasoning is as follows: assume that we have a large number of random candidates (the $\mathbf{t}_{unif}^{(i)}$) and one point close to the lattice L (the point \mathbf{e}), then Heuristic 2 says that we can always distinguish \mathbf{e} from the candidates (since it has maximum value of g_W). The contradiction comes from the fact that in reality, for T large enough, many of candidates will be closer to L than \mathbf{e} and therefore no algorithm can distinguish them [18]. This gives rise to what [20] calls the “contradictory regime” where an algorithm would somehow be able to distinguish indistinguishable distributions.

We first compare this regime to that of our algorithm and we then discuss the statistical model chosen by [20] in Heuristic 1.

⁷ We overload the notation GH: in [20], $\text{GH}(m)$ corresponds to our $\text{GH}(L)$ for L of volume 1, that is $\text{vol}(B_m)^{-1/m}$.

6.1 Almost Complementary Regimes

In Sect. 4.3, we have applied our main theorem to a concrete instance and derived that⁸ for a typical LWE problem where the ratio m/n is fixed (and not too close to 0 or 1), q is large and the error follows a discrete Gaussian of parameter σ_e , our algorithm works as soon as

$$q^{1-n/m} \geq (\frac{1}{s} + 2\sigma_e)\sqrt{e} \tag{24}$$

where

$$N = \frac{\text{poly}(m) + n \log_2(q)}{\delta^2}, \quad \delta = \frac{1}{10}e^{-ms^2\sigma_e^2/2}.$$

In our attack, T is the number of guesses that the algorithm makes, that is $T = q^{n_{\text{guess}}}$. In order to match [20, page 21], we will choose s so that $\ln T = N\epsilon^2$:

$$\begin{aligned} \ln T = N\epsilon^2 &\Leftrightarrow n_{\text{guess}} \ln(q) = \frac{\text{poly}(m) + n \log_2(q)}{\delta^2} \epsilon^2 \\ &\Leftrightarrow n_{\text{guess}} \ln(q) = (\text{poly}(m) + n \log_2(q)) 100e^{2ms^2\sigma_e^2/2} e^{-2\pi\sigma_e^2\ell^2} \\ &\Leftrightarrow \frac{n_{\text{guess}} \ln(q)}{100(\text{poly}(m) + n \log_2(q))} = e^{(ms^2 - 2\pi\ell^2)\sigma_e^2}. \end{aligned}$$

Note that $n_{\text{guess}} < n < m$ so for large enough value of m , the left-hand side of this expression is smaller than 1 (recall that $\text{poly}(m)$ comes from the choice of N so we can always make it slightly bigger to artificially increase the denominator if we want). It follows that we can always choose s such that $\ln T = N\epsilon^2$ in such a way that (24) holds (see [39, Appendix G.2]) and therefore Theorem 5 ensures that our algorithm works in this regime.

We will now compare this with [20]’s contradictory regime. This regime, defined in [20, page 21] is when⁹

$$r \text{GH}(L_q(\mathbf{A}_{\text{dual}})) < \sqrt{\frac{m}{2\pi}}\sigma_e, \quad \text{where } r = T^{-1/m}. \tag{25}$$

Note here that the lattice is \mathbf{A}_{dual} because [20] modularizes the algorithm by separating the lattice in which dual-distinguishing is done, with the part of the lattice that is enumerated over (see Sect. 6.2). Indeed, this regime comes from Heuristic 1 and the lattice in question is the one where dual vectors are generated.

Recall that for the algorithm to work, \mathbf{A} and therefore \mathbf{A}_{dual} must have full rank, so $\det(L_q(\mathbf{A}_{\text{dual}})) = q^{m-n_{\text{dual}}}$. Now observe that

$$\frac{r \text{GH}(L_q(\mathbf{A}_{\text{dual}}))}{\sqrt{\frac{m}{2\pi}}\sigma_e} = \frac{T^{-1/m} \sqrt{\frac{m}{2\pi e}} q^{1-n_{\text{dual}}/m}}{\sqrt{\frac{m}{2\pi}}\sigma_e} = \frac{q^{-n_{\text{guess}}/m} q^{1-n_{\text{dual}}/m}}{\sqrt{e}\sigma_e}.$$

⁸ Under some mild technical simplification to make the computation easier.

⁹ Recall that because of the difference between the normal distribution and the discrete Gaussian, we have $\sigma = \sigma_e/\sqrt{2\pi}$ in our analysis, see [39, Appendix G.1].

Recall that $n = n_{\text{dual}} + n_{\text{guess}}$ so the contradictory regimes corresponds to

$$q^{1-n/m} < \sigma_e \sqrt{e}. \quad (26)$$

Comparing between the working regime (24) and the contradictory one (26), and recalling that we can choose s as large as we want, we observe that they do not overlap and the bounds only differ by a factor of two. This suggest that, for our algorithm, the “theoretically working” regime and the contradictory regime almost characterize whether the dual attack will work or not. However, the next section will explain that those regimes are based on different distributions of targets.

6.2 On the Distribution of Targets

The authors of [20] decided to modularize the algorithm by separating the lattice in which dual-distinguishing is done ($L_q(\mathbf{A}_{\text{dual}})$) from the part of the lattice that is enumerated over ($L_q(\mathbf{A}_{\text{guess}})$). In fact, Heuristic 1 only mentions the dual-distinguishing and not the enumeration. This however, poses a difficulty because it is clear that the “targets” ($\mathbf{b} - \mathbf{A}_{\text{guess}} \tilde{\mathbf{s}}_{\text{guess}}$ in our terminology, $\mathbf{t}_{\text{unif}}^{(i)}$ in Heuristic 1) are not arbitrary but have some structure.

The authors of [20] decided to model the statistics of the targets in a way that is independent of the actual choice of $\mathbf{A}_{\text{guess}}$: they chose the uniform distribution over the fundamental domain of $L_q(\mathbf{A}_{\text{dual}})$. In the case of [32] and our algorithm, the algorithm exclusively works over integers which is why we propose Heuristic 2 as an integer-version of Heuristic 1. This means that we now have two different settings:

- In Heuristic 2, $\mathbf{t}_{\text{unif}}^{(i)}$ is sampled uniformly in \mathbb{Z}^m/L .
- In reality, $\mathbf{t}_{\text{unif}}^{(i)} = \mathbf{e} + \mathbf{x}^{(i)}$ where $\mathbf{x}^{(i)}$ can be any vector in $L' \setminus q\mathbb{Z}^m$ where L' is another random q -ary lattice, chosen independently of L but fixed in the algorithm. In our algorithm, $L = L_q(\mathbf{A}_{\text{dual}})$ and $L' = L_q(\mathbf{A}_{\text{guess}})$.

Indeed, a key point in the proof of Theorem 5 is to show that points of the form $\mathbf{e} + \mathbf{x}^{(i)}$ as described are always far away from L , a fact that does not hold for completely uniform targets. As a result, with high probability over the choice of \mathbf{A} , the targets (except for the correct guess) are *all bounded away* from 0 in the dual lattice. For uniform targets, the argument of [20] is statistical in nature: while there can be very short vectors, they are unlikely and the contradiction comes from the fact that if we try too many targets, we will eventually find a short one and get a false-positive. On the other hand, our algorithm and analysis is not statistical: for the vast majority of choices of \mathbf{A} , all targets satisfy the bound unconditionally and we can safely look at all targets without the risk of any false-positive.

In conclusion of this section, it seems that the contradictory regime of [20] nicely complements the working regime of our algorithm. On the other hand, the statistical model that underlines this contradictory regime and what happens in our algorithm are different. We leave it as an open question to explain exactly why the two regimes seem to align perfectly.

7 Open Questions

We have analysed formally a dual attack in the spirit of [32]. However, as noted in [20], the algorithm used by [32] produces many short dual vectors *in a sublattice* L'' of $L_q^\perp(\mathbf{A}_{\text{dual}})$ (instead of the entire $L_q^\perp(\mathbf{A}_{\text{dual}})$). In other words, W is roughly the set of vectors of L'' in a ball and therefore g_W does not exactly measure the distance to L but rather to a more complicated lattice. This fact makes the analysis of g_W considerably more challenging and we believe that more research is needed to understand how this affects the choice of the parameters.

Another issue that we have avoided is that of modulus switching. Indeed, while [32] claims that this techniques bring significant improvements in the complexity, [20] claims that geometric arguments contradicts this statement. We leave as an open problem the study of a modification of our algorithm that would include modulus switching. We believe that a formal analysis would be the best way to resolve this issue. A priori, we do not see any major reason why this could not be analysed formally but it may prove to be a nontrivial technical challenge due to the effects of rounding modulo p on the uniform distribution modulo q . We note in this direction that the approach of [16] of using lattice codes instead of modulus switching might be a better fit for a formal analysis.

Finally, we have analyzed the case where the algorithm has access to m LWE samples in dimension n , and our algorithm typically requires $m \approx 2n$ to have a good complexity. In practice, however, it is common to only have n samples, something that our algorithm cannot handle. While there is a standard technique to deal with this, namely sampling in the lattice

$$\{(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^m \times \mathbb{Z}^{n_{\text{dual}}} : \mathbf{x}^T \mathbf{A}_{\text{dual}} = \mathbf{y} \bmod q\},$$

we leave it as future work to include this improvement to our analysis.

Acknowledgments. We thank the anonymous TCC reviewers for pointing out an error in a previous version of this paper where we misunderstood the contradictory regime of [20]. We thank Martin Albrecht for his helpful comments on a previous version of this paper. We thank Léo Ducas for helpful discussions on the statistical model of [20]. We thank the anonymous EUROCRYPT reviewers for valuable comments and suggestions. Y.S. is supported by EPSRC grant EP/W02778X/2.

References

1. Lattice Attacks on NTRU and LWE: A History of Refinements, p. 15-40. London Mathematical Society Lecture Note Series, Cambridge University Press (2021)
2. Aggarwal, D., Chen, Y., Kumar, R., Shen, Y.: Improved classical and quantum algorithms for the shortest vector problem via bounded distance decoding (2020). <https://arxiv.org/abs/2002.07955>
3. Aggarwal, D., Chen, Y., Kumar, R., Shen, Y.: Improved (provable) algorithms for the shortest vector problem via bounded distance decoding. In: Bläser, M., Monmege, B. (eds.) 38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, Saarbrücken, Germany, 16–19 March 2021 (Virtual

- Conference). LIPIcs, vol. 187, pp. 4:1–4:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2021). <https://doi.org/10.4230/LIPIcs.STACS.2021.4>
4. Aggarwal, D., Dadush, D., Regev, O., Stephens-Davidowitz, N.: Solving the shortest vector problem in 2^n time using discrete gaussian sampling: Extended abstract. In: Servedio, R.A., Rubinfeld, R. (eds.) Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, 14–17 June 2015, pp. 733–742. ACM (2015). <https://doi.org/10.1145/2746539.2746606>
 5. Aharonov, D., Regev, O.: Lattice problems in $np \cap comp$. J. ACM **52**(5), 749–765 (2005). <https://doi.org/10.1145/1089023.1089025>
 6. Ajtai, M.: The shortest vector problem in L2 is NP-hard for randomized reductions (extended abstract). In: Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing, STOC 1998, pp. 10–19. Association for Computing Machinery, New York (1998). <https://doi.org/10.1145/276698.276705>
 7. Albrecht, M.R.: On dual lattice attacks against small-secret LWE and parameter choices in HELIB and SEAL. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10211, pp. 103–129. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56614-6_4
 8. Albrecht, M.R., Gheorghiu, V., Postlethwaite, E.W., Schanck, J.M.: Estimating quantum speedups for lattice sieves. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020. LNCS, vol. 12492, pp. 583–613. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64834-3_20
 9. 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>
 10. Albrecht, M.R., Shen, Y.: Quantum augmented dual attack. Cryptology ePrint Archive, Paper 2022/656 (2022). <https://eprint.iacr.org/2022/656>
 11. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange: a new hope. In: Proceedings of the 25th USENIX Conference on Security Symposium, SEC 2016, pp. 327–343. USENIX Association, USA (2016)
 12. Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. Math. Ann. **296**, 625–635 (1993)
 13. 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, 10–12 January 2016, pp. 10–24. SIAM (2016). <https://doi.org/10.1137/1.9781611974331.ch2>
 14. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, STOC 2013, pp. 575–584. Association for Computing Machinery, New York (2013). <https://doi.org/10.1145/2488608.2488680>
 15. Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-LWE and security for key dependent messages. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 505–524. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22792-9_29
 16. Carrier, K., Shen, Y., Tillich, J.P.: Faster dual lattice attacks by using coding theory (2022). <https://eprint.iacr.org/2022/1750>
 17. Dadush, D., Regev, O., Stephens-Davidowitz, N.: On the closest vector problem with a distance guarantee. In: 2014 IEEE 29th Conference on Computational Complexity (CCC), pp. 98–109 (2014). <https://doi.org/10.1109/CCC.2014.18>

18. Debris, T., Ducas, L., Resch, N., Tillich, J.P.: Smoothing codes and lattices: systematic study and new bounds. *Cryptology ePrint Archive*, Paper 2022/615 (2022). <https://eprint.iacr.org/2022/615>
19. Debris-Alazard, T., Remaud, M., Tillich, J.P.: Quantum reduction of finding short code vectors to the decoding problem. *Cryptology ePrint Archive*, Paper 2021/752 (2021). <https://eprint.iacr.org/2021/752>
20. Ducas, L., Pulles, L.N.: Does the dual-sieve attack on learning with errors even work? *IACR Cryptol. ePrint Arch.*, p. 302 (2023). <https://eprint.iacr.org/2023/302>
21. Dürr, C., Høyer, P.: A quantum algorithm for finding the minimum. CoRR **quant-ph/9607014** (1996). <http://arxiv.org/abs/quant-ph/9607014>
22. Erez, U., Litsyn, S., Zamir, R.: Lattices which are good for (almost) everything. *IEEE Trans. Inf. Theory* **51**(10), 3401–3416 (2005). <https://doi.org/10.1109/TIT.2005.855591>
23. Espitau, T., Joux, A., Kharchenko, N.: On a dual/hybrid approach to small secret LWE. In: Bhargavan, K., Oswald, E., Prabhakaran, M. (eds.) *INDOCRYPT 2020*. LNCS, vol. 12578, pp. 440–462. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-65277-7_20
24. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, pp. 197–206. STOC 2008. Association for Computing Machinery, New York (2008). <https://doi.org/10.1145/1374376.1374407>
25. Guo, Q., Johansson, T.: Faster dual lattice attacks for solving LWE with applications to CRYSTALS. In: Tibouchi, M., Wang, H. (eds.) *ASIACRYPT 2021*. LNCS, vol. 13093, pp. 33–62. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-92068-5_2
26. Høyer, P., Mosca, M., de Wolf, R.: Quantum search on bounded-error inputs. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) *ICALP 2003*. LNCS, vol. 2719, pp. 291–299. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-45061-0_25
27. Jaques, S., Rattew, A.G.: QRAM: a survey and critique (2023)
28. Klein, P.: Finding the closest lattice vector when it’s unusually close. In: *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2000*, pp. 937–941. Society for Industrial and Applied Mathematics, USA (2000)
29. Kuperberg, G.: A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM J. Comput.* **35**(1), 170–188 (2005). <https://doi.org/10.1137/S0097539703436345>
30. Loeliger, H.A.: Averaging bounds for lattices and linear codes. *IEEE Trans. Inf. Theory* **43**(6), 1767–1773 (1997). <https://doi.org/10.1109/18.641543>
31. Macbeath, A.M., Rogers, C.A.: Siegel’s mean value theorem in the geometry of numbers. *Math. Proc. Cambridge Philos. Soc.* **54**(2), 139–151 (1958). <https://doi.org/10.1017/S0305004100033302>
32. MATZOV: Report on the Security of LWE: Improved Dual Lattice Attack, April 2022. <https://doi.org/10.5281/zenodo.6412487>
33. Meyer-Hilfinger, C., Tillich, J.P.: Rigorous foundations for dual attacks in coding theory. *Cryptology ePrint Archive*, Paper 2023/1460, Accepted to TCC (2023). <https://eprint.iacr.org/2023/1460>
34. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on gaussian measures. In: *45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 372–381 (2004). <https://doi.org/10.1109/FOCS.2004.72>

35. Micciancio, D., Regev, O.: Lattice-Based Cryptography, pp. 147–191. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-540-88702-7_5
36. Micciancio, D., Voulgaris, P.: Faster exponential time algorithms for the shortest vector problem. In: Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, pp. 1468–1480. Society for Industrial and Applied Mathematics, USA (2010)
37. Montanaro, A.: Quantum speedup of Monte Carlo methods. Proc. Royal Soc. Math. Phys. Eng. Sci. **471**(2181), 20150301 (2015). <https://doi.org/10.1098/rspa.2015.0301>
38. Nguyen, P.Q., Vidick, T.: Sieve algorithms for the shortest vector problem are practical. J. Math. Cryptol. **2**(2), 181–207 (2008). <https://doi.org/10.1515/JMC.2008.009>
39. Pouly, A., Shen, Y.: Provable dual attacks on learning with errors. Cryptology ePrint Archive, Paper 2023/1508 (2023). <https://eprint.iacr.org/2023/1508>
40. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, STOC 2005, pp. 84–93. Association for Computing Machinery, New York (2005). <https://doi.org/10.1145/1060590.1060603>
41. Rogers, C.A.: Mean values over the space of lattices. Acta Math. **94**, 249–287 (1955). <https://doi.org/10.1007/BF02392493>
42. Schnorr, C.P.: Lattice reduction by random sampling and birthday methods. In: Alt, H., Habib, M. (eds.) STACS 2003. LNCS, vol. 2607, pp. 145–156. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-36494-3_14
43. Schnorr, C.: A hierarchy of polynomial time lattice basis reduction algorithms. Theor. Comput. Sci. **53**(2), 201–224 (1987)
44. Siegel, C.L.: A mean value theorem in geometry of numbers. Ann. Math. **46**(2), 340 (1945). <https://doi.org/10.2307/1969027>. <https://www.jstor.org/stable/1969027?origin=crossref>
45. Stephens-Davidowitz, N.: On the Gaussian measure over lattices. Ph.d. thesis, New York University (2017)
46. 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>
47. Zamir, R., Nazer, B., Kochman, Y., Bistriz, I.: Lattice Coding for Signals and Networks: A Structured Coding Approach to Quantization, Modulation and Multiuser Information Theory. Cambridge University Press (2014). <https://doi.org/10.1017/CBO9781139045520>