Improved Estimation of Key Enumeration with Applications to Solving LWE

Alessandro Budroni[1] and Erik Mårtensson[2,3]
Cryptography Research Center, Technology Innovation Institute, Abu Dhabi, UAE[1]
alessandro.budroni@tii.ae
Selmer Center, Department of Informatics, University of Bergen, Norway[2]
Department of Electrical and Information Technology, Lund University, Sweden[3]
erik.martensson@{uib.no,eit.lth.se}

***Abstract*—In post-quantum cryptography (PQC), Learning With Errors (LWE) is one of the dominant underlying mathematical problems. For example, in NIST's PQC standardization process, the Key Encapsulation Mechanism (KEM) protocol chosen for standardization was Kyber, an LWE-based scheme. Recently the dual attack surpassed the primal attack in terms of concrete complexity for solving the underlying LWE problem for multiple cryptographic schemes, including Kyber. The dual attack consists of a reduction part and a distinguishing part. When estimating the cost of the distinguishing part, one has to estimate the expected cost of enumerating over a certain number of positions of the secret key. Our contribution consists of giving a polynomial-time approach for calculating the expected complexity of such an enumeration procedure. This allows us to revise the complexity of the dual attack on the LWE-based protocols Kyber, Saber and TFHE. For all these schemes we improve upon the total bit-complexity in both the classical and the quantum setting.**

As our method of calculating the expected cost of enumeration is fairly general, it might be of independent interest in other areas of cryptography or even in other research areas.

I. Introduction

Introduced by Regev in 2005 [1], the Learning With Errors Problem (LWE) is a computational problem that has been used as a building block for several quantum-resistant cryptographic primitives. A consistent number of schemes in each round of NIST's Post-Quantum Standardization Process [2] base their security on the hardness of LWE. One of them is Kyber, which was chosen as the standard algorithm for encryption. Saber is another LWE-based scheme, which is very similar to Kyber and made it to the third round of the competition. It is also possible to build Fully Homomorphic Encryption (FHE) on LWE. TFHE is such an encryption scheme, based on [3].

Cryptanalysis of LWE is an active area of research that encompasses various techniques, including combinatorial methods like the Blum-Kalai-Wasserman (BKW) algorithm [4], algebraic methods [5], and lattice-reduction-based approaches, such as the primal attack [6] and the (recently) most successful dual attack [7]–[10]. Both BKW and the dual attack, in their most recent variants, include a subroutine consisting of enumerating a vector with entries from a non-uniform distribution. Previous works dealt with this problem either using unexplained models for estimating the cost of enumeration or using unnecessarily pessimistic upper limit formulas [9], [10].

The contribution of this manuscript is to provide a new and more accurate method to estimate the cost of such an enumeration procedure. Our key realization is that the frequencies of the different possible secret coefficient values follow a multinomial distribution, meaning that the number of unique probabilities for different possible keys is only polynomial in the number of positions we enumerate over. This allows us to precisely calculate the expected cost of key enumeration in polynomial time.

We integrate this new method into the complexity estimation of the dual attack and obtain new security estimates for the widely studied lattice-based schemes Kyber, Saber and TFHE, both for the classic and quantum case scenario. Furthermore, our contribution is general enough that it easily can be applied to any situation where enumeration over a vector sampled from a non-uniform distribution is needed.

The remaining part of the paper is organized as follows. In Section II, we present notations and necessary background. In Section III we introduce our new key enumeration approach, while in Section IV we apply it to some lattice-based protocols. Finally, in Section V we give the conclusions.

II. Preliminaries

A. Notation

We denote the set of the integer, rational and real numbers with $\mathbb{Z}, \mathbb{Q}, \mathbb{R}$ respectively. For a positive integer p, we write $\mathbb{Z}_p = \mathbb{Z}/p\mathbb{Z}$. Upper case letters, e.g. M, denote matrices, and bold lower case letters, e.g. $\boldsymbol{v}$, represent column vectors. We represent with v_j the j-th component of $\boldsymbol{v}$. We let $\log(\cdot)$ denote the 2-logarithm. The notation $\|\boldsymbol{v}\|$ denotes the Euclidean norm of $\boldsymbol{v} \in \mathbb{R}^n$ defined as

$$\|\boldsymbol{v}\| = \sqrt{v_1^2 + \cdots + v_n^2}.$$

For a discrete distribution X, its entropy is defined as

$$H(X) := -\mathbb{E}(\log(X)) = -\sum_k p(x_k) \cdot \log(p(x_k)). \quad (1)$$

B. Quantum Search Algorithms

Grover's algorithm is a way of efficiently searching for elements in an unstructured set. Let $\mathcal{S}$ be a finite set of N

objects of which $t \leq N$ are *targets*. An oracle $\mathcal{O}$ identifies the targets if, for every $s \in \mathcal{S}$, $\mathcal{O}(s) = 1$ if s is a target, $\mathcal{O}(s) = 0$ otherwise. Classically, one needs $O(N/t)$ queries to the oracle to identify a target. Grover provided a quantum algorithm that identifies a target with only $O(\sqrt{N/t})$ queries to the oracle [11]. *Amplitude amplification* is a subsequent work that generalizes Grover's search algorithm [12].

C. Lattices and Reduction Algorithms

A **lattice** is a discrete additive subgroup of $\mathbb{R}^n$. Let $B = \{\boldsymbol{b}_1, ..., \boldsymbol{b}_m\} \in \mathbb{R}^n$ be a set of linearly independent vectors. We define the lattice generated by B as

$$\mathcal{L}(B) = \mathcal{L}(\boldsymbol{b}_1, ..., \boldsymbol{b}_m) = \left\{ \boldsymbol{v} \in \mathbb{R}^n : \boldsymbol{v} = \sum_{i=1}^{m} \alpha_i \mathbf{b}_i,\ \alpha_i \in \mathbb{Z} \right\}.$$

Unless differently specified, we will consider full-rank lattices, i.e. $n = m$.

Typically, lattice reduction algorithms such as LLL or BKZ [13]–[15], take as input a basis B of the lattice and return another basis with short and nearly orthogonal vectors. Lattice sieving consists of a class of algorithms, initiated with the work of Ajtai et al. [16], to solve the Shortest Vector Problem (SVP). These are usually used internally by BKZ as an SVP oracle . They allow us to compute a large number of short vectors and they have an estimated complexity of $2^{c\beta+o(\beta)}$, where β is the dimension of the lattice and c is a constant equal to 0.292 for classical computers [17]. This constant can be improved quantumly to 0.2653 using Grover's algorithm [18]. It was recently further improved to 0.2570 by using more sophisticated quantum methods [19][1].

D. Learning With Errors and Gaussian Distributions

Definition 1: Let n be a positive integer, q a prime and χ_s, χ_e two probability distributions over $\mathbb{Z}$. Fix a secret vector $\boldsymbol{s} \in \mathbb{Z}^n$ whose entries are sampled according to χ_s. Denote by $\mathcal{A}_{\mathbf{s},\chi_e}$ the probability distribution on $\mathbb{Z}_q^n \times \mathbb{Z}_q$ obtained by sampling $\boldsymbol{a} \in \mathbb{Z}_q^n$ uniformly at random, sampling an error $e \in \mathbb{Z}$ from χ_e and returning

$$(\boldsymbol{a}, z) = (\boldsymbol{a}, \langle \boldsymbol{a} \cdot \boldsymbol{s} \rangle + e \mod q) \in \mathbb{Z}_q^n \times \mathbb{Z}_q.$$

- The *search* Learning With Errors (LWE) problem is to find the secret vector $\boldsymbol{s}$ given a fixed number of samples from $\mathcal{A}_{\mathbf{s},\chi_e}$.
- The *decision* Learning With Errors (LWE) problem is to distinguish between samples drawn from $\mathcal{A}_{\mathbf{s},\chi_e}$ and samples drawn uniformly from $\mathbb{Z}_q^n \times \mathbb{Z}_q$.

Consider m LWE samples

$$(\mathbf{a}_1, z_1), (\mathbf{a}_2, z_2), \ldots, (\mathbf{a}_m, z_m) \leftarrow \mathcal{A}_{s,\chi}.$$

Then, one can represent such an LWE instance in a matrix-vector form as

$$(A, \mathbf{z}) = (A, A\mathbf{s} + \mathbf{e} \mod q) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$$

[1]or even 0.2563 according to a recent preprint [20].

where A is an $m \times n$ matrix with rows $\mathbf{a}_1^T, \mathbf{a}_2^T, \ldots, \mathbf{a}_m^T$, $\mathbf{z} = (z_1, z_2, \ldots, z_m)$, and $\mathbf{e}$ is the vector of errors $(e_1, e_2, \ldots, e_m)$.

In theory, one usually instantiates χ_s and χ_e as the *discrete Gaussian distribution* $D_{\mathbb{Z},\sigma}$ defined as the probability distribution that to each $a \in \mathbb{Z}$ assigns the probability

$$\frac{\rho_{0,\sigma}(a)}{\sum_{d \in \mathbb{Z}} \rho_{0,\sigma}(d)} = \frac{\exp\left(-\pi a^2/2\sigma^2\right)}{\sum_{d \in \mathbb{Z}} \exp(-\pi d^2/2\sigma^2)},$$

where $\rho_{0,\sigma}(x)$ is the probability distribution function of the Gaussian distribution $N(0, \sigma)$ with mean 0 and variance σ^2. In practice, it is more common to use a centered Binomial distribution $\mathbf{B}_\eta$, which takes values in $[-\eta, \eta]$ or a uniform distribution $\mathcal{U}\{a, b\}$, which takes values in $[a, b]$.

Given an LWE problem instance, there exists a polynomial-time transformation [21], [22] that makes the secret vector follow the same distribution as the error's distribution χ_e.

E. Distinguishing Attacks to LWE

a) Dual Attack: The first attack on LWE performed on the so-called *dual* lattice was introduced in [7]. While the earlier versions of this attack were efficient only for instances with very small coefficients (e.g. $\mathbf{s} \in \{-1, 0, 1\}^n$), thanks to some recent contributions [8]–[10], the attack now also applies to secrets with not-so-small coefficients.

Let $(A, \mathbf{b} = A\mathbf{s} + \mathbf{e} \mod q)$ be an $m \times n$ LWE instance, for $m \geq n$ where the secret $\mathbf{s}$ and the error $\mathbf{e}$ have been sampled from a discrete normal distribution with mean zero and standard deviations σ_s and σ_e respectively. Partition the matrix A as $(A_1 \parallel A_2)$ and, in correspondence, the secret $\mathbf{s}$ as $(\mathbf{s}_1 \parallel \mathbf{s}_2)$. Consider the following pair

$$(A_2, \mathbf{b} - A_1\tilde{\mathbf{s}}_1 \mod q). \tag{2}$$

For $\tilde{\mathbf{s}}_1 = \mathbf{s}_1$ we have that

$$\mathbf{b} - A_1\tilde{\mathbf{s}}_1 = A_2\mathbf{s}_2 + \mathbf{e} \mod q$$

and therefore (2) is a new LWE instance with reduced dimension. If $\tilde{\mathbf{s}}_1 \neq \mathbf{s}_1$, then (2) is uniform.

By enumerating over all possible vectors $\tilde{\mathbf{s}}_1$ of $\mathbf{s}_1$, one can distinguish the right guess as follows. Let $\mathcal{R}$ be an algorithm (e.g. BKZ, lattice sieving) that returns pairs $(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^{m \times n}$ such that $\mathbf{y}^T = (\mathbf{y}_1 \parallel \mathbf{y}_2)^T = \mathbf{x}^T A \mod q$, and $\mathbf{x}$ and $\mathbf{y}_2$ are *short*. Then, for $\tilde{\mathbf{s}}_1 = \mathbf{s}_1$, we have that

$$\mathbf{x}^T(\mathbf{b} - A_1\mathbf{s}_1) = \mathbf{x}^T(A_2\mathbf{s}_2 + \mathbf{e}) = \mathbf{y}_2^T\mathbf{s}_2 + \mathbf{x}^T\mathbf{e}. \tag{3}$$

This quantity is distributed approximately according to a discrete normal distribution with mean zero and variance $\|\mathbf{x}\|^2\sigma_s^2 + \|\mathbf{y}_2\|^2\sigma_e^2$. The choice for reduction algorithm $\mathcal{R}$ determines the expected length of the vectors $\mathbf{x}$ and $\mathbf{y}_2$, and therefore, the ability to distinguish (3) from uniform random.

b) BKW algorithm: In its original development, the Blum-Kalai-Wasserman (BKW) algorithm was proposed as a subexponential algorithm for solving the Learning Parity with Noise (LPN) problem [23]. Later, it has been applied to LWE [4], and further developed with new ideas such as Lazy-Modulus-Switching, Coded-BKW and smooth-Lazy-Modulus-Switching [24]–[28].

The BKW algorithm can be seen as a variant of the dual attack where the reduction is performed using combinatorial methods instead of lattice reduction. For this reason, techniques and improvements developed for BKW on the distinguishing stage have been successfully applied to the dual attack too. More generally, the BKW algorithm has the disadvantage of requiring an exponential number of samples ($m \gg n$) to perform reduction. On the other hand, one typically has more control over the outcome (i.e. reduced samples).

III. Improved Estimation of Key Enumeration

Consider the problem of guessing the random value X sampled from a discrete probability distribution with mass function $p_k := P(X = x_k)$. Without loss of generality, we assume it to be *non-increasing* (i.e. $p_0 \geq p_1 \geq p_2 \geq \dots$). The optimal strategy is obviously to guess that $X = x_0$, followed by guessing that $X = x_1$, and so on. The expected number of guesses until the right value is found with this strategy is

$$G(X) = \sum_i i \cdot p_i. \tag{4}$$

Massey showed in [29] that

$$G(X) \geq \frac{1}{4} 2^{H(X)}.$$

He also showed why there is no such formula for upper limiting $G(X)$ in terms of $H(X)$.

Now consider a sample of n values, each one drawn independently from the same distribution with mass function $(p_0, p_1, \dots, p_{r-1})$. When enumerating all the possible samples of $\mathbf{s}$ on these n positions, we want to do so in decreasing order of probability until we find the solution. Since the total number of outcomes is equal to r^n, we cannot simply compute the probability of every single outcome, sort all the probabilities and then compute the expectation that way. However, we can use the fact that the frequencies of each possible secret value follow the multinomial distribution [30]. The number of outcomes where k_0 values are equal to x_0, k_1 values are equal to x_1 and so on until k_{r-1} values are equal to x_{r-1}, where $\sum_{i=0}^{r-1} k_i = n$, is

$$\binom{n}{k_0, \dots, k_{r-1}} = \frac{n!}{k_0! k_1! \cdots k_{r-1}!}. \tag{5}$$

Notice that all these outcomes have exactly the same probability of

$$\prod_{l=0}^{r-1} p_l^{k_l}. \tag{6}$$

The total number of unique probabilities is only

$$\mu = \binom{n+r-1}{n} = \frac{(n+r-1)\cdots(n+1)}{(r-1)!} = \frac{(n+r-1)!}{(r-1)!n!}. \tag{7}$$

For a fixed number r this expression is $\mathcal{O}(n^{r-1})$. Thus, for a sparse distribution the number of unique probabilities is low enough to be computed and sorted efficiently (i.e. in polynomial time w.r.t. n).

Denote the unique probabilities by $p'_0, p'_1, \dots, p'_{\mu-1}$, such that $p'_0 \geq p'_1 \geq \cdots \geq p'_{\mu-1}$. Let f_i denote the number of times p'_i occurs. Also let $F_i = \sum_{j=0}^{i-1} f_j$. Now we can express the expected number of guesses to make until we find the right one from (4), as

$$\sum_{i=0}^{\mu-1} p'_i \left(F_i + \sum_{j=1}^{f_i} j \right) = \sum_{i=0}^{\mu-1} p'_i \left(F_i + \frac{f_i(f_i+1)}{2} \right). \tag{8}$$

Since (8) has $\mathcal{O}(n^{r-1})$ terms and each term can be computed efficiently, the whole expression can be computed efficiently for small values of r.

A. Quantum setting

Consider again random values sampled from a discrete probability with mass function $(p_0, p_1, p_2, ...)$. With a quantum computer, as shown in [31] using amplitude amplification as a tool, the expected number of guesses to find the right value is

$$G_{\text{qc}}(X) = \sum_i \sqrt{i} \cdot p_i. \tag{9}$$

Using the Cauchy-Schwartz inequality we have that

$$\begin{aligned} G_{\text{qc}}(X) &= \sum_i \sqrt{i \cdot p_i} \cdot \sqrt{p_i} \leq \sqrt{\sum_i i \cdot p_i \cdot \sum_i p_i} \\ &= \sqrt{\sum_i i \cdot p_i} = \sqrt{G(X)}. \end{aligned} \tag{10}$$

Here, our method for computing the estimated cost of the enumeration of (9) still applies, with a minor twist. In this setting (8) changes to

$$\sum_{i=0}^{\mu-1} p'_i \left(\sum_{j=1}^{f_i} \sqrt{F_i + j} \right). \tag{11}$$

We can rewrite $\sum_{j=1}^{f_i} \sqrt{F_i + j} = \sum_{j=1}^{F_i+f_i} \sqrt{j} - \sum_{j=1}^{F_i} \sqrt{j}$. Now, to compute (11) efficiently we only need to have an efficient and precise formula for computing $f(n) = \sum_{i=1}^{n} \sqrt{i}$. For $n \leq 30$ we can pre-compute the expression. For $n > 30$ using the Euler-Maclaurin formula [32], we can derive the function

$$f(n) \approx \zeta(-0.5) + \frac{1}{2} n^{\frac{1}{2}} + \frac{2}{3} n^{\frac{3}{2}} + \frac{1}{24} n^{-\frac{1}{2}} - \frac{1}{1920} n^{-\frac{5}{2}} + \frac{1}{9216} n^{-\frac{9}{2}}, \tag{12}$$

where $\zeta(\cdot)$ is the Riemann zeta function, which approximates the sum with a relative error that is smaller than or equal to machine epsilon.

B. Further optimizations

If for two outcomes x_1 and x_2 we have $P(x_1) = P(x_2)$, then we can merge these terms to speed up the calculation of the enumeration.

Also, more generally, if throughout the enumeration we have two lists of values $[x_1, x_2, \ldots, x_k]$ and $[x'_1, x'_2, \ldots, x'_k]$ and $P([x_1, x_2, \ldots, x_k]) = P([x'_1, x'_2, \ldots, x'_k])$, then we can also merge these two terms.

IV. Application to Lattice-based Schemes

In the Matzov version of the dual attack on LWE, the n positions of the secret $\mathbf{s}$ are divided up into three parts, k_{lat}, k_{fft} and k_{enum}. The attack first performs lattice reduction on k_{lat} positions. In the second phase it enumerates, in decreasing order of probability, all possible secrets on k_{enum} positions. For each such secret it performs an FFT on k_{fft} positions and checks if it has found the correct solution. Rewriting [9, Theorem 5.1] asymptotically we get the following formula for the cost of the distinguishing part of the dual attack.

$$\mathcal{O}\big(G(\chi^{k_{\text{enum}}}) \cdot (D + p^{k_{\text{fft}}})\big), \tag{13}$$

where D is the number of samples needed to distinguish the secret and $\chi^{k_{\text{enum}}}$ refers to the distribution of k_{enum} values sampled independently from the distribution χ. The fact that the cost is additive in D and $p^{k_{\text{fft}}}$ means that it is best to keep these two terms of similar size. Quantumly however, the cost is proportional to the square root of the number of samples needed to distinguish the secret, the cost of enumeration and the cost of performing the FFT quantumly [10, (4)]. More concretely the cost is

$$\mathcal{O}\Big(\sqrt{D} \cdot p^{k_{\text{fft}}/2} \cdot G_{\text{qc}}(\chi^{k_{\text{enum}}}) \cdot \text{poly}(\log(n))\Big). \tag{14}$$

The drastically reduced cost of distinguishing is the main source of the quantum improvement that [10] achieves compared to [9]. Notice the more than quadratic speed-up of $G_{\text{qc}}(\chi^{k_{\text{enum}}})$ over $G(\chi^{k_{\text{enum}}})$, as shown in (10). In practice this speed-up means that it is optimal for the schemes studied in this paper to do enumeration only and let $k_{\text{fft}} = 0$.

In Matzov [9], it was assumed that the expected cost of enumerating over k_{enum} positions is $2^{k_{\text{enum}} \cdot H(\chi)}$, without any explanation. In [10], this problem was addressed. They developed an upper limit formula for the expected cost of enumerating over k_{enum} positions sampled from a Discrete Gaussian distribution with a specified standard deviation σ. When estimating the expected cost of enumerating over the secret of an actual scheme, they simply approximated the secret distribution as a Discrete Gaussian with the same standard deviation, see Table III. In the quantum setting they developed a similar model.

Using the method detailed in Section III, in both the classical and quantum setting we can calculate the expected cost of enumeration numerically with arbitrarily good precision, to compare against the models of [9], [10]. Since all the schemes use sparse (and symmetric) distributions for the secret, our method is very efficient at computing the expectations.

A classical comparison is illustrated in Figure 1, for the expected cost of enumeration for Kyber512/FireSaber. The exhaustive cost is the obvious upper limit of guessing every possible key. Notice that while the Matzov numbers are a bit too optimistic, they are actually closer to the exact numbers than the Albrecht/Shen model is. Notice that the gaps between the different models increase with the dimension.

Figure 2 covers the quantum setting. Notice that there is a consistent gap between the expected cost according to the Albrecht/Shen model and the exact value, which increases very slowly with the number of dimensions.

Table I shows the state-of-the-art of solving the underlying LWE problem using the dual attack for the different schemes and models considered in [10]. We briefly summarize the models here. The models CC, CN and C0 are increasingly optimistic models for the cost of the dual attack on classical computers. GE19 refers to the most pessimistic quantum model from [33]. QN and Q0 correspond to CN and C0, but with the classical lattice sieving of [17] replaced by the quantum lattice sieving of [19]. Finally, QN [10] and Q0 [10] refer to the works of [10], where quantum speed-ups of the FFT and the enumeration are applied. All the numbers are computed using the script from [10].

Table II shows the updated state-of-the-art. These are achieved by replacing Albrecht's and Shen's upper limit formulas for enumeration by the exact values, as described in Section III[2]. For all schemes and all models we show improvements, but the magnitude of the improvements vary. Our largest improvements are for the TFHE schemes, where the secret follows a uniform distribution, meaning that a Discrete Gaussian is a particularly bad approximation.

Very recently, another preprint of an improved version of the dual attack of Matzov was published [34]. There they introduce a modified way of enumerating over the secret. Compared to the results from [10] they achieve comparable levels of improvements to us, in the classical setting. They enumerate over the secret in a different way, meaning that our improved estimate of the cost of enumeration does not apply in their setting. However, they do not provide a quantum version of their improved algorithm, the setting where our contribution is the most impactful.

TABLE I
PREVIOUS STATE-OF-THE-ART.

Scheme	CC	CN	C0	GE19	QN	Q0	Q0 [10]	QN [10]
Kyber512	139.2	134.4	115.4	139.5	124.4	102.7	119.3	99.7
Kyber768	196.1	190.6	173.7	191.9	175.3	154.6	168.3	150.0
Kyber1024	262.4	256.1	241.8	252.0	234.5	215.0	225.6	208.4
LightSaber	138.5	133.1	113.7	138.4	122.7	101.1	118.9	98.9
Saber	201.4	195.9	179.2	196.2	179.9	159.4	173.8	155.0
FireSaber	263.5	258.2	243.8	253.1	235.9	216.7	228.1	210.8
TFHE630	118.2	113.3	93.0	120.2	105.2	83.0	100.8	80.7
TFHE1024	122.0	117.2	95.4	123.9	108.5	84.8	105.6	83.2

[2]The code for computing these numbers will be made available.

TABLE II
UPDATED STATE-OF-THE-ART.

Scheme	CC	CN	C0	GE19	QN	Q0	Q0 [10]	QN [10]
Kyber 512	138.7	133.8	115.0	139.1	123.6	102.4	118.0	98.4
Kyber 768	194.8	190.0	172.9	190.6	174.5	154.5	166.3	148.0
Kyber 1024	260.6	254.5	240.6	251.0	233.4	214.5	223.2	206.2
LightSaber	137.5	132.6	113.3	138.0	122.3	101.0	117.6	97.7
Saber	200.9	195.6	178.5	196.1	179.3	159.2	172.4	153.8
FireSaber	262.9	256.9	242.6	252.8	235.3	216.4	226.2	208.8
TFHE630	115.7	111.3	92.1	118.2	103.9	82.8	95.6	76.8
TFHE1024	120.4	115.6	94.8	122.8	107.7	84.5	101.7	80.4

TABLE III
THE SECRET DISTRIBUTION AND ITS STANDARD DEVIATION, FOR EACH SCHEME.

Scheme	Distribution	Standard deviation
Kyber512	$\mathbf{B}_3$	$\sqrt{6}/2$
Kyber768	$\mathbf{B}_2$	1
Kyber1024	$\mathbf{B}_2$	1
LightSaber	$\mathbf{B}_5$	$\sqrt{10}/2$
Saber	$\mathbf{B}_4$	$\sqrt{8}/2$
FireSaber	$\mathbf{B}_3$	$\sqrt{6}/2$
TFHE630	$\mathcal{U}\{0,1\}$	$1/2$
TFHE1024	$\mathcal{U}\{0,1\}$	$1/2$

A. Applications to BKW

As discussed in Section II-E, the techniques introduced in Section III apply to the BKW algorithm too. In the setting of [27], [28], the secret coefficients are discrete Gaussian with a relatively large standard deviation, taken from the distributions of the LWE Darmstadt Challenges [35]. The authors perform enumeration over all possible secret values within 3 standard deviations for each position. By instead enumerating over the secret coefficients in decreasing order of probability, one would see improvements similar to those of the dual attack.

V. CONCLUSIONS

The method presented in this paper improves over previous estimations for key-enumeration used in the literature. As a direct application, we used it to revise the state-of-the-art complexities for the dual attack against Kyber, Saber and TFHE. Future research directions include the application of this methods to other areas in cryptanalysis where an enumeration of a non-uniform vector is required. Furthermore, thanks to its generality, the method might find application also in areas outside the context of cryptography.

Acknowledgments. We thank Qian Guo, Martin Albrecht and Yixin Shen for helpful discussions on the topic. We also thank the anonymous reviewers of ISIT23 for useful comments on the initial submission of this paper. Erik Mårtensson was supported by the project "Kvantesikker Kryptografi" from the National Security Authority of Norway.

REFERENCES

[1] O. Regev, "On Lattices, Learning with Errors, Random Linear Codes, and Cryptography," in *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*. ACM, 2005, pp. 84–93.

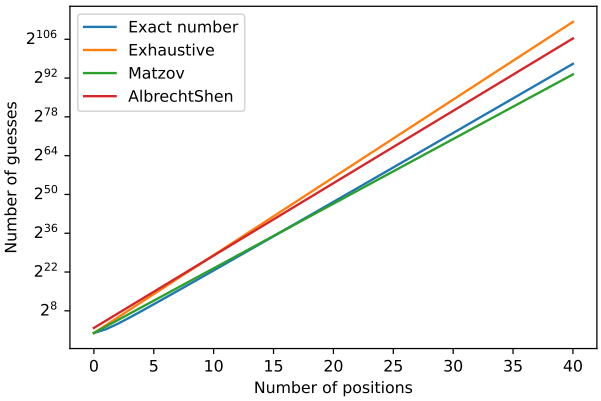

Fig. 1. The expected cost of enumeration in the classic setting.

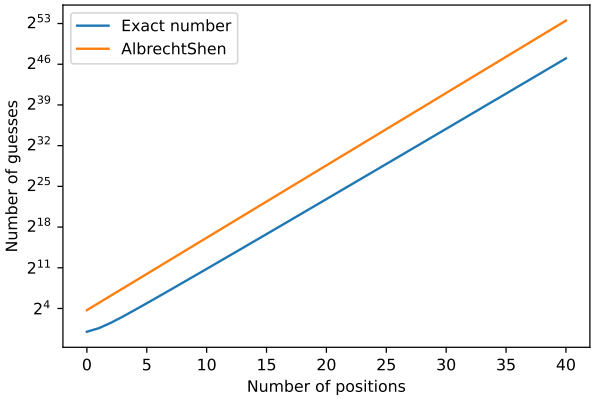

Fig. 2. The expected cost of enumeration in the quantum setting.

[2] NIST, "Post-quantum cryptography standardization," https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization.

[3] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, "Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds," in *Advances in Cryptology – ASIACRYPT 2016*, J. H. Cheon and T. Takagi, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 3–33.

[4] M. R. Albrecht, C. Cid, J. Faugère, R. Fitzpatrick, and L. Perret, "On the complexity of the BKW algorithm on LWE," *Des. Codes Cryptogr.*, vol. 74, no. 2, pp. 325–354, 2015.

[5] S. Arora and R. Ge, "New Algorithms for Learning in Presence of Errors," in *Automata, Languages and Programming*, ser. LNCS, vol. 6755. Springer, 2011, pp. 403–415.

[6] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe, "Post-Quantum Key Exchange: A New Hope," in *Proceedings of the 25th USENIX Conference on Security Symposium*, ser. SEC'16. USENIX, 2016, p. 327–343.

[7] D. Micciancio and O. Regev, *Lattice-based Cryptography*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 147–191. [Online]. Available: https://doi.org/10.1007/978-3-540-88702-7_5

[8] Q. Guo and T. Johansson, "Faster dual lattice attacks for solving LWE with applications to CRYSTALS," in *Advances in Cryptology – ASIACRYPT 2021, Part IV*, ser. Lecture Notes in Computer Science, M. Tibouchi and H. Wang, Eds., vol. 13093. Singapore: Springer, Heidelberg, Germany, Dec. 6–10, 2021, pp. 33–62.

[9] MATZOV, "Report on the Security of LWE: Improved Dual Lattice Attack," Apr. 2022. [Online]. Available: https://doi.org/10.5281/zenodo.6412487

[10] M. R. Albrecht and Y. Shen, "Quantum augmented dual attack," Cryptology ePrint Archive, Paper 2022/656, 2022, https://eprint.iacr.org/2022/656. [Online]. Available: https://eprint.iacr.org/2022/656

[11] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, ser. STOC '96. New York, NY, USA: Association for Computing Machinery, 1996, p. 212–219. [Online]. Available: https://doi.org/10.1145/237814.237866

[12] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, "Quantum amplitude amplification and estimation," *Contemporary Mathematics*, vol. 305, pp. 53–74, 2002.

[13] A. K. Lenstra, H. W. Lenstra, and L. Lovasz, "Factoring Polynomials with Rational Coefficients," *Mathematische Annalen*, vol. 261, pp. 515–534, 1982.

[14] C. Schnorr, "A hierarchy of polynomial time lattice basis reduction algorithms," *Theor. Comput. Sci.*, vol. 53, no. 2, p. 201–224, jun 1987.

[15] Y. Chen and P. Q. Nguyen, "BKZ 2.0: Better Lattice Security Estimates," in *Advances in Cryptology – ASIACRYPT 2011*, ser. LNCS, vol. 7073. Springer, 2011, pp. 1–20.

[16] M. Ajtai, R. Kumar, and D. Sivakumar, "A sieve algorithm for the shortest lattice vector problem," in *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, ser. STOC '01. New York, NY, USA: Association for Computing Machinery, 2001, p. 601–610. [Online]. Available: https://doi.org/10.1145/380752.380857

[17] A. Becker, L. Ducas, N. Gama, and T. Laarhoven, "New Directions in Nearest Neighbor Searching with Applications to Lattice Sieving," in *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2016, pp. 10–24.

[18] T. Laarhoven, M. Mosca, and J. van de Pol, "Finding shortest lattice vectors faster using quantum search," *Designs, Codes, and Cryptography*, vol. 77, pp. 375 – 400, 2015.

[19] A. Chailloux and J. Loyer, "Lattice sieving via quantum random walks," in *Advances in Cryptology – ASIACRYPT 2021*, M. Tibouchi and H. Wang, Eds. Cham: Springer International Publishing, 2021, pp. 63–91.

[20] X. Bonnetain, A. Chailloux, A. Schrottenloher, and Y. Shen, "Finding many collisions via reusable quantum walks," Cryptology ePrint Archive, Paper 2022/676, 2022, https://eprint.iacr.org/2022/676. [Online]. Available: https://eprint.iacr.org/2022/676

[21] B. Applebaum, D. Cash, C. Peikert, and A. Sahai, "Fast cryptographic primitives and circular-secure encryption based on hard learning problems," in *Advances in Cryptology – CRYPTO 2009*, ser. Lecture Notes in Computer Science, S. Halevi, Ed., vol. 5677. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 16–20, 2009, pp. 595–618.

[22] P. Kirchner, "Improved generalized birthday attack," Cryptology ePrint Archive, Report 2011/377, 2011, https://eprint.iacr.org/2011/377.

[23] A. Blum, A. T. Kalai, and H. Wasserman, "Noise-tolerant learning, the parity problem, and the statistical query model," in *Symposium on the Theory of Computing*, 2000.

[24] M. R. Albrecht, J.-C. Faugère, R. Fitzpatrick, and L. Perret, "Lazy modulus switching for the BKW algorithm on LWE," in *PKC 2014: 17th International Conference on Theory and Practice of Public Key Cryptography*, ser. Lecture Notes in Computer Science, H. Krawczyk, Ed., vol. 8383. Buenos Aires, Argentina: Springer, Heidelberg, Germany, Mar. 26–28, 2014, pp. 429–445.

[25] Q. Guo, T. Johansson, and P. Stankovski, "Coded-BKW: Solving LWE using lattice codes," in *Advances in Cryptology – CRYPTO 2015, Part I*, ser. Lecture Notes in Computer Science, R. Gennaro and M. J. B. Robshaw, Eds., vol. 9215. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 16–20, 2015, pp. 23–42.

[26] P. Kirchner and P.-A. Fouque, "An improved BKW algorithm for LWE with applications to cryptography and lattices," in *Advances in Cryptology – CRYPTO 2015, Part I*, ser. Lecture Notes in Computer Science, R. Gennaro and M. J. B. Robshaw, Eds., vol. 9215. Santa Barbara, CA, USA: Springer, Heidelberg, Germany, Aug. 16–20, 2015, pp. 43–62.

[27] A. Budroni, Q. Guo, T. Johansson, E. Mårtensson, and P. S. Wagner, "Making the bkw algorithm practical for lwe," in *Progress in Cryptology – INDOCRYPT 2020*, K. Bhargavan, E. Oswald, and M. Prabhakaran, Eds. Cham: Springer International Publishing, 2020, pp. 417–439.

[28] A. Budroni, Q. Guo, T. Johansson, E. Mårtensson, and P. S. Wagner, "Improvements on making bkw practical for solving lwe," *Cryptography*, vol. 5, no. 4, 2021. [Online]. Available: https://www.mdpi.com/2410-387X/5/4/31

[29] J. Massey, "Guessing and entropy," in *Proceedings of 1994 IEEE International Symposium on Information Theory*, 1994, pp. 204–.

[30] Wikipedia contributors, "Multinomial distribution — Wikipedia, the free encyclopedia," 2023, accessed 2023-01-24. [Online]. Available: https://en.wikipedia.org/wiki/Euler-Maclaurin_formula

[31] A. Montanaro, "Quantum search with advice," in *Theory of Quantum Computation, Communication, and Cryptography*, W. van Dam, V. M. Kendon, and S. Severini, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 77–93.

[32] Wikipedia contributors, "Euler–Maclaurin formula — Wikipedia, the free encyclopedia," 2023, accessed 2023-01-10. [Online]. Available: https://en.wikipedia.org/wiki/Euler-Maclaurin_formula

[33] C. Gidney and M. Ekerå, "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits," *Quantum*, vol. 5, p. 433, Apr. 2021. [Online]. Available: https://doi.org/10.22331/q-2021-04-15-433

[34] K. Carrier, Y. Shen, and J.-P. Tillich, "Faster dual lattice attacks by using coding theory," Cryptology ePrint Archive, Paper 2022/1750, 2022, https://eprint.iacr.org/2022/1750.

[35] "TU Darmstadt Learning with Errors Challenge," https://www.latticechallenge.org/lwe_challenge/challenge.php, accessed: 2023-01-24.