

# spKEX: An optimized lattice-based key-exchange

Sauvik Bhattacharya<sup>1</sup>, Oscar Garcia-Morchon<sup>1</sup>, Ronald Rietman<sup>1</sup>, and Ludo Tolhuizen<sup>1</sup>

<sup>1</sup>Philips Research, Eindhoven, The Netherlands  
{sauvik.bhattacharya, oscar.garcia-morchon, ronald.rietman,  
ludo.tolhuizen}@philips.com

## Abstract

The advent of large-scale quantum computers has resulted in significant interest in quantum-safe cryptographic primitives. Lattice-based cryptography is one of the most attractive post-quantum cryptographic families due to its well-understood security, efficient operation and versatility. However, LWE-based schemes are still relatively bulky and slow.

In this work, we present spKEX, a forward-secret, post-quantum, unauthenticated lattice-based key-exchange scheme that combines four techniques to optimize performance. spKEX relies on Learning with Rounding (LWR) to reduce bandwidth; it uses sparse and ternary secrets to speed up computations and reduce failure probability; it applies an improved key reconciliation scheme to reduce bandwidth and failure probability; and computes the public matrix  $A$  by means of a permutation to improve performance while allowing for a fresh  $A$  in each key exchange.

For a quantum security level of 128 bits, our scheme requires 30% lesser bandwidth than the LWE-based key-exchange proposal Frodo [9] and allows for a fast implementation of the key exchange.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Notation</b>	<b>4</b>
<b>3</b>	<b>The spKEX Key-Exchange Protocol</b>	<b>5</b>
3.1	Underlying Problem . . . . .	5
3.2	Protocol Description . . . . .	6
<b>4</b>	<b>Analysis of Operational Correctness</b>	<b>7</b>
4.1	Difference between Raw Keys . . . . .	8
4.2	Reconciliation Condition . . . . .	8
<b>5</b>	<b>Security Analysis</b>	<b>9</b>
5.1	Attacks based on Lattice Reduction . . . . .	9
5.2	Hybrid Attack . . . . .	12
5.3	Choice of Hamming Weight . . . . .	13
5.4	Unbiased Keys . . . . .	15
5.5	Choice and Computation of the Public Parameter $\mathbf{A}$ . . . . .	15
5.6	Parameter Selection . . . . .	16
<b>6</b>	<b>Performance Analysis</b>	<b>17</b>
6.1	Bandwidth performance . . . . .	17
6.2	Computational Performance . . . . .	18
<b>7</b>	<b>Conclusions &amp; Future Work</b>	<b>19</b>
<b>8</b>	<b>Acknowledgments</b>	<b>19</b>
<b>A</b>	<b>Upper bound on Failure Probability</b>	<b>22</b>
<b>B</b>	<b>Distinguishing Advantage in Dual Attack</b>	<b>24</b>

# 1 Introduction

The exchange of sensitive information, e.g. financial, military or private data, over communication systems requires solutions to ensure the confidentiality of these information transactions. Confidentiality can be enabled by means of cryptographic primitives such as public-key encryption [24] and key-establishment [15]. However, the approach of quantum computing, in combination with Shor’s [26] and Grover’s [17] algorithms has resulted in a need for quantum-resistant algorithms since a majority of the world’s communication today depend on cryptography that is vulnerable to cryptanalysis by quantum computers. This need is recognized by organizations such as NIST and ETSI that are currently standardizing such solutions.

*Key-exchange* is one of the most sensitive and urgent application areas where post-quantum cryptography is required. The reason is that an attacker who collects encrypted data of interest to him (even assuming that forward-secret cryptography is used) will be able to decrypt this data when a quantum computer becomes available. Thus, this paper focuses on this primitive.

Lattice-based cryptography is a promising candidate for quantum-resistant cryptography due to its (relatively) good performance, versatility in different cryptographic schemes and resistance against all known quantum algorithms. In particular, the Learning with Errors (LWE) problem [23] is a hard mathematical problem with quantum reductions to the worst-case hard lattice problems GAPSVP and SIVP [23] and classical reductions to GAPSVP [21, 10]. In one of the flavors of LWE, the attacker is given many pairs  $(a_i, \{b_i = a_i s + e_i \pmod{q}\})$  and his task is to recover  $s$ , where  $a_i$  and  $s$  are randomly chosen vectors from a uniform distribution and  $e_i$  are randomly taken from a Gaussian distribution. The Learning with Rounding (LWR) problem [6] is a deterministic variant of the LWE problem that replaces standard Gaussian errors with errors introduced by rounding to a smaller modulus. The motivation is to achieve higher efficiency due to the difficulty of sampling from a Gaussian distribution and to reduce the ciphertext length [11]. For both LWE and LWR, *small-secret* and *sparse* variants are also possible, i.e., ones in which the secret is sampled from a binary or ternary distribution, secrets that are sparse, and secrets that are both sparse and small [13].

The schemes of Ding *et al* [16] and the Frodo scheme due to Bos *et al* [9] are post-quantum, unauthenticated key-exchange schemes based on the Learning with Errors (LWE) problem. They use public keys derived using secret vectors and Gaussian noise to establish a shared secret between two entities in an ephemeral manner. The sampling from a Gaussian distribution while creating the LWE error leads to a significant overhead. The Lizard scheme of Cheon *et al* [14] is a public-key encryption (PKE) scheme that uses the reduction of the LWE problem to the Learning with Rounding (LWR) problem [6, 8, 5] in order to replace the slower Gaussian sampling with more efficient rounding for incorporating noise. Furthermore, Lizard proposes the use of sparse, ternary secrets in order to increase efficiency, and is based on both LWE and the LWR problem with sparse-ternary secrets (sparse-ternary LWR or SP-TERLWR). However,

these schemes are still relatively bulky and the open question is if they can be further optimized for real-world usage scenarios.

This document describes an unauthenticated, ephemeral lattice-based key-exchange – SPKEX, based on the LWR problem with sparse-ternary secrets. This key-exchange proposal builds on the lattice-based KEX proposed by [16] and [9]. Performance improvements are achieved by a combination of four design techniques:

- Introduction of noise by using rounding [6], which is computationally more efficient than sampling noise from a discrete Gaussian distribution [2, 9, 14] and allows reducing bandwidth requirements.
- Use of the generalized reconciliation mechanism from [27] to achieve a (sufficiently low) target failure probability with smaller configuration parameters.
- The application of sparse ternary secrets to optimize computations.
- Introduction of a computationally inexpensive approach to refresh the public parameter  $\mathbf{A}$  by permuting a fixed, master public parameter. This aims at preventing pre-computation or backdoor-like attacks [22] while being more efficient than purely random approaches [9, 2].

Although some of the techniques above are known, they have not been used in actual key exchanges yet. For instance, Frodo mentions the potential usage of LWR but it does not describe the design of such a key exchange nor does it quantify its security level. Thus, the present work aims at showing how efficient a lattice-based key exchange can be made if all known techniques that improve performance are applied together, while providing a thorough analysis of the corresponding security and operational results. The remainder of this paper is organized as follows: Section 2 describes our notation. Section 3 describes SPKEX. Section 4, 5, and 6 respectively analyze the failure probability, security, and performance of SPKEX. Section 7 concludes this paper.

## 2 Notation

A vector is denoted in small letters, bold-face, as  $\mathbf{v}$ , and is always considered as a column vector. Matrices are denoted in capital letters, bold-face, as  $\mathbf{M}$ . We denote the inner product of two vector  $\mathbf{v}_1$  and  $\mathbf{v}_2$  of equal length as  $(\mathbf{v}_1, \mathbf{v}_2)$ ; that is,  $(\mathbf{v}_1, \mathbf{v}_2) = \mathbf{v}_1^\top \cdot \mathbf{v}_2$ . The norm of a vector is the Euclidean or  $\ell_2$  norm unless mentioned otherwise. Logarithms are assumed to be base 2 unless otherwise specified. We denote by  $\|\mathbf{M}\|_\infty$  the maximum of the absolute value of the entries of  $\mathbf{M}$ . That is, for any  $n \times m$  matrix  $\mathbf{M}$ ,

$$\|\mathbf{M}\|_\infty = \max\{|M_{i,j}| \mid 1 \leq i \leq n, 1 \leq j \leq m\}.$$

For each real number  $x$ , we write  $\lfloor x \rfloor$  for  $x$  rounded downwards to the closest integer, so

$$\lfloor x \rfloor = \max\{n \in \mathbb{Z} \mid n \leq x\}.$$

Moreover,  $\lfloor x \rfloor$  denotes the integer closest to  $x$  (rounded upwards in case of a tie), so  $\lfloor x \rfloor = \lfloor x + \frac{1}{2} \rfloor$ . For each integer  $x$  and each integer positive integer  $p$ , we define  $\langle x \rangle_p$  and  $\{x\}_p$  as the integer in  $(0, p-1]$  and in  $(-p/2, p/2]$ , respectively, that is equivalent to  $x$  modulo  $p$ . Note that for each integer  $x$  and each positive integer  $p$ ,  $\{x\}_p = \langle x + \lfloor \frac{p}{2} \rfloor \rangle_p - \lfloor \frac{p}{2} \rfloor$ . We extend these notations to matrices by component wise application.

If  $\chi$  is a probability distribution over a set  $S$ , then  $x \stackrel{\$}{\leftarrow} \chi$  denotes sampling  $x \in S$  according to  $\chi$ . If  $S$  is a set, then  $U(S)$  denotes the uniform distribution on  $S$ , and we denote sampling of  $x$  uniformly and randomly from  $S$  either by  $x \stackrel{\$}{\leftarrow} U(S)$  or by  $x \stackrel{\$}{\leftarrow} S$ . For a matrix  $\mathbf{M}$  or vector  $\mathbf{v}$ , the notation  $\mathbf{M} \stackrel{\$}{\leftarrow} \chi$  and  $\mathbf{v} \stackrel{\$}{\leftarrow} \chi$  signifies that the entries of  $\mathbf{M}$  and  $\mathbf{v}$  are drawn independently according to distribution  $\chi$ . For a real  $\sigma > 0$ , we use  $D_\sigma$  to denote the discrete Gaussian distribution, which has support  $\mathbb{Z}$  and assigns a probability proportional to  $\exp(-\pi x^2/\sigma^2)$  to each  $x \in \mathbb{Z}$ . For  $n, h \in \mathbb{Z}$  and  $0 \leq h \leq n$ ,  $\mathcal{HWT}_n(h)$  represents the set of vectors from  $\{-1, 0, 1\}^n$  of Hamming weight  $h$ , i.e., each such vector contains exactly  $h$  non-zero entries.

### 3 The spKEX Key-Exchange Protocol

#### 3.1 Underlying Problem

The security of our key-exchange scheme is based on the hardness of the sparse-ternary LWR problem:

**Definition 1 (Sparse-Ternary LWR (sp-terLWR))** *Let  $n, p, q, h$  be positive integers.*

*The **search-sp-terLWR** is to find  $\mathbf{s} \in \{-1, 0, 1\}^n$  of Hamming weight  $h$  from arbitrarily many independent samples of the form  $(\mathbf{a}_i, b_i = \lfloor \frac{p}{q} \cdot \langle \mathbf{a}, \mathbf{s} \rangle_q \rfloor)$ .*

*The **decision-sp-terLWR** problem is to distinguish  $(\mathbf{a}, b = \lfloor \frac{p}{q} \cdot \langle \mathbf{a}, \mathbf{s} \rangle_q \rfloor)$  from the uniform distribution over  $\mathbb{Z}_q^n \times \mathbb{Z}_p$  with non-negligible advantage, for a fixed  $\mathbf{s} \in \{-1, 0, 1\}^n$  of Hamming weight  $h$ .*

According to [14], the hardness of this problem can be obtained from that of LWE with similar secret distributions since the reduction from LWE to LWR is independent of the secret's distribution [8]. Even if a formal reduction is not presented in [14], this statement is sound. Consider LWE with uniform noise on  $V = (-\frac{q}{2p}, \frac{q}{2p}] \cap \mathbb{Z}$ . For a given secret  $\mathbf{s}$ , the attacker observes  $(\mathbf{a}, b) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$  where  $b = \langle \mathbf{a}, \mathbf{s} \rangle_q + e \pmod{q}$ , and  $e \stackrel{\$}{\leftarrow} V$ . We thus have that  $\frac{p}{q} \cdot b = \frac{p}{q} \cdot \langle \mathbf{a}, \mathbf{s} \rangle_q + \frac{p}{q} \cdot e \pmod{p}$ . As  $\frac{p}{q} \cdot e \in (-1/2, 1/2]$ , we infer that if  $b$  is a multiple of  $\frac{q}{p}$ , then  $\frac{p}{q} \cdot b = \lfloor \frac{p}{q} \cdot \langle \mathbf{a}, \mathbf{s} \rangle_q \rfloor$ . Hence, if  $\mathbf{s}$  can be obtained from  $m$  LWR samples, then it can be obtained from  $\frac{q}{p}m$  samples from LWE with uniform noise on  $(-\frac{q}{2p}, \frac{q}{2p}]$ . This is a similar reasoning as in [8, Thm.3]. Note that this reduction is valid for all distributions of the components of  $\mathbf{s}$ .

Results on the hardness of LWR are relevant for spKEX. It is shown in [8] that the search and decision LWR problems are at least as hard as the corre-

sponding LWE problems when the number of problem instances  $m$  is bounded. This however requires an increase of the LWR problem dimension by a factor of  $\log q$ . In [5] this reduction is tightened, showing that an LWR problem is at least as hard as an LWE problem with uniform, bounded noise that supplies a greater number of samples but has the *same* dimension  $n$  and modulus  $q$ . This reduction is notable since it is the first to preserve the dimension  $n$  and modulus  $q$  between LWE (albeit with uniform, bounded noise) and LWR without resorting to noise-flooding [6] to re-normalize the rounding errors. In spKEX, an attacker has a bounded number of samples since it is a (ephemeral) key exchange in which parameters are refreshed. Thus, LWR and its variants are suitable choices for the construction of an ephemeral key exchange protocol, considering the bound on  $m$ .

### 3.2 Protocol Description

This section describes our key-exchange protocol between Alice and Bob that is sketched in Table 1. The protocol begins with the choice of the following public parameters:

- $q$ : The large modulus of the LWR problem.
- $p$ : The rounding modulus of the LWR problem, an integer multiple of  $2^{B+b_h+1}$  and satisfying  $p|q$ .
- $n$ : Dimension of the LWR problem, also representing the dimension of the public matrix  $\mathbf{A} \in \mathbb{Z}_q^{n \times n}$ .
- $h$ : The Hamming weight of secrets.
- $\bar{n}, \bar{m}$ : Number of LWR problem instances created by Alice and Bob, respectively.
- $B$ : Number of bits extracted per coefficient of each party's raw key.
- $b_h$ : The number of reconciliation bits per coefficient of each party's raw key.
- $\mathbf{A}_{\text{master}}$ : The LWR master public matrix from which a fresh instance is generated for each key-exchange session by means of a random fresh permutation  $\Pi$ .
- A hash function `hash`.

Based on these public system parameters, Alice and Bob use the spKEX protocol detailed in Table 1 to agree on a shared secret  $\mathbf{K}$ , and its hash – the final symmetric key  $\mu$ . This protocol is similar to [9] with a number of key differences.

*First*, the public matrix  $\mathbf{A} \in \mathbb{Z}_q^{n \times n}$  is generated from a fixed, master public matrix  $\mathbf{A}_{\text{master}}$  and a random permutation  $\Pi$  as  $\mathbf{A} = \Pi(\mathbf{A}_{\text{master}})$ . This allows

Table 1: spKEX Protocol

Alice	Bob
1. Choose a random permutation $\Pi$ .	
2. $\mathbf{A} = \Pi(\mathbf{A}_{\text{master}}) \in \mathbb{Z}_q^{n \times n}$ .	
3. $\mathbf{S}_a \in \{0, \pm 1\}^{n \times \bar{n}} \leftarrow \mathcal{HWT}_n(h)^{\bar{n}}$ .	
4. $\mathbf{P}_a = \langle \lfloor (\frac{p}{q}) \langle \mathbf{A} \cdot \mathbf{S}_a \rangle_q \rfloor \rangle_p \in \mathbb{Z}_p^{n \times \bar{n}}$ .	
	5. $\Pi, \mathbf{P}_a$ $\rightarrow$
	6. $\mathbf{A} = \Pi(\mathbf{A}_{\text{master}}) \in \mathbb{Z}_q^{n \times n}$ .
	7. $\mathbf{S}_b \in \{0, \pm 1\}^{n \times \bar{m}} \leftarrow \mathcal{HWT}_n(h)^{\bar{m}}$ .
	8. $\mathbf{P}_b = \langle \lfloor (\frac{p}{q}) \langle \mathbf{S}_b^T \cdot \mathbf{A} \rangle_q \rfloor \rangle_p \in \mathbb{Z}_p^{\bar{m} \times n}$ .
	9. $\mathbf{K}_b = \langle \mathbf{S}_b^T \cdot \mathbf{P}_a \rangle_p \in \mathbb{Z}_p^{\bar{m} \times \bar{n}}$ .
	10. $\mathbf{H} = \lfloor \frac{2^{B+b_h}}{p} \langle \mathbf{K}_b \rangle_p / 2^B \rfloor \in \mathbb{Z}_{2^{b_h}}^{\bar{m} \times \bar{n}}$ .
	11. $\mathbf{P}_b, \mathbf{H}$ $\leftarrow$
12. $\mathbf{K}_a = \langle \mathbf{P}_b \cdot \mathbf{S}_a \rangle_p \in \mathbb{Z}_p^{\bar{m} \times \bar{n}}$ .	
13. $\mathbf{K} = \langle \lfloor \frac{2^B}{p} \mathbf{K}_a - \frac{1}{2^{b_h}} \mathbf{H} - \frac{1}{2^{b_h+1}} + \frac{1}{2} \rfloor \rangle_{2^B}$ .	13. $\mathbf{K} = \lfloor \frac{2^B}{p} \mathbf{K}_b \rfloor$
14. $\mu = \text{hash}(\mathbf{K})$	14. $\mu = \text{hash}(\mathbf{K})$

efficient computation every time of a fresh  $\mathbf{A}$ , i.e., this approach is a trade-off between the performance advantages of a fixed matrix  $\mathbf{A}$  [22] and the security considerations of having a fresh random  $\mathbf{A}$  to prevent pre-computation or backdoor-like attacks [9], e.g., lattice-reduction.

*Second*, spKEX uses sparse ternary secrets  $\mathbf{S}_a, \mathbf{S}_b$  of which each column has Hamming weight  $h$  so that operations are more efficient. Because  $\mathbf{S}_a$  and  $\mathbf{S}_b$  are fresh for each session, spKEX achieves forward secrecy. *Third*, the public-keys are computed by applying rounding. This avoids expensive sampling of Gaussian noise, and reduces bandwidth requirements to exchange public-keys. Furthermore, the “raw keys”  $\mathbf{K}_a$  and  $\mathbf{K}_b$  can be obtained by performing computations modulo the (smaller) modulus  $p$ .

*Fourth*, to achieve an *exact* key agreement we use the reconciliation mechanism from [27] that allows for larger differences between the raw keys than the mechanism from [22].

## 4 Analysis of Operational Correctness

In this section, we analyze the operational correctness of spKEX. That is, we show that for a proper choice of parameters, Alice and Bob agree on a common key with a high probability.

## 4.1 Difference between Raw Keys

We define the  $n \times \bar{n}$  matrix  $\mathbf{E}_a$  and the  $\bar{m} \times n$  matrix  $\mathbf{E}_b$  as

$$\mathbf{E}_a = \mathbf{P}_a - \frac{p}{q} \langle \mathbf{A} \cdot \mathbf{S}_a \rangle_q = \lfloor \frac{p}{q} \langle \mathbf{A} \cdot \mathbf{S}_a \rangle_q \rfloor - \frac{p}{q} \langle \mathbf{A} \cdot \mathbf{S}_a \rangle_q \quad (1)$$

$$\mathbf{E}_b = \mathbf{P}_b - \frac{p}{q} \langle \mathbf{S}_b \cdot \mathbf{A} \rangle_q = \lfloor \frac{p}{q} \langle \mathbf{S}_b \cdot \mathbf{A} \rangle_q \rfloor - \frac{p}{q} \langle \mathbf{S}_b \cdot \mathbf{A} \rangle_q \quad (2)$$

The definition of the  $\lfloor \cdot \rfloor$  function implies that

$$\|\mathbf{E}_a\|_\infty \leq \frac{1}{2} \text{ and } \|\mathbf{E}_b\|_\infty \leq \frac{1}{2}. \quad (3)$$

In fact, all entries of  $\mathbf{E}_a$  and  $\mathbf{E}_b$  are integer multiples of  $\frac{p}{q}$ . From the definitions of  $\mathbf{K}_a, \mathbf{K}_b, \mathbf{E}_a$  and  $\mathbf{E}_b$ , it follows that

$$\begin{aligned} \mathbf{K}_a &\equiv \mathbf{P}_b \cdot \mathbf{S}_a \\ &= \left\{ \frac{p}{q} \cdot (\mathbf{S}_b^\top \cdot \mathbf{A} + q \cdot \mathbf{U}_1) + \mathbf{E}_b \right\} \cdot \mathbf{S}_a \quad \text{where } \mathbf{U}_1 \in \mathbb{Z}^{\bar{m} \times n} \\ &\equiv \frac{p}{q} (\mathbf{S}_b^\top \cdot \mathbf{A} \cdot \mathbf{S}_a) + \mathbf{E}_b \cdot \mathbf{S}_a \pmod{p}, \end{aligned} \quad (4)$$

and similarly

$$\mathbf{K}_b \equiv \frac{p}{q} (\mathbf{S}_b^\top \cdot \mathbf{A} \cdot \mathbf{S}_a) + \mathbf{S}_b^\top \cdot \mathbf{E}_a \pmod{p}. \quad (5)$$

As a consequence,

$$\mathbf{K}_a - \mathbf{K}_b \equiv \mathbf{E}_b \cdot \mathbf{S}_a - \mathbf{S}_b^\top \cdot \mathbf{E}_a \pmod{p}. \quad (6)$$

## 4.2 Reconciliation Condition

As shown in [27], Alice and Bob agree on entry  $i, j$  of  $\mathbf{K}$  if

$$\{(\mathbf{K}_a - \mathbf{K}_b)_{i,j}\}_p \leq \Delta := \frac{p}{2^{B+1}} - \frac{p}{2^{B+b_h+1}}. \quad (7)$$

We invoke (6) to analyze  $\{(\mathbf{K}_a - \mathbf{K}_b)\}_p$ . Each column of  $\mathbf{S}_a$  and each row of  $\mathbf{S}_b^\top$  has  $h$  non-zero entries, and each non-zero entry equals 1 or  $-1$ . Moreover, all entries of  $\mathbf{E}_a$  and  $\mathbf{E}_b$  are in the interval  $(-1/2, 1/2]$ . As a consequence, all entries of  $\mathbf{E}_b \cdot \mathbf{S}_a$  and  $\mathbf{S}_b^\top \cdot \mathbf{E}_a$  are in  $(-\frac{1}{2}h, \frac{1}{2}h]$ , and so  $\|\mathbf{E}_b \cdot \mathbf{S}_a - \mathbf{S}_b^\top \cdot \mathbf{E}_a\|_\infty \leq h$ . Hence, if  $h \leq \Delta$ , then Alice and Bob always agree on a common key.

We wish to have a more relaxed condition that ensures that Alice and Bob agree on a common key *with high probability*. In Appendix A we show the following result. If we model the entries of each row of  $\mathbf{E}_a$  and each row of  $\mathbf{E}_b$  as being independently and uniformly distributed on  $(-1/2, 1/2]$ , then

$$\Pr(\text{Alice and Bob do not agree on a common key}) \leq 4 \cdot \bar{m} \cdot \bar{n} \cdot 2^{h \cdot f((\Delta+1)/h)}, \quad (8)$$

$$\text{where } f(a) = \frac{1}{\ln 2} \min\left\{\ln\left(\frac{e^t - e^{-t}}{2t}\right) - at \mid t > 0\right\}. \quad (9)$$

We have not been able to derive a closed formula for  $f$ , but the minimization can easily be performed numerically. We defer details to Appendix A.

## 5 Security Analysis

Our security analysis considers four aspects, namely, *a*) attacks using lattice reduction and parameter estimations, *b*) specialized attacks exploiting sparse ternary secrets, *c*) biases in the computed keys, and *d*) security considerations regarding the computation of the public parameter  $\mathbf{A}$ .

### 5.1 Attacks based on Lattice Reduction

The use of ephemeral secrets and refreshed public parameters in the key-exchange strictly limits the number of samples available to an attacker, and restricts relevant attacks to those based on lattice reduction – the *primal* or decoding attack [4] and the *dual* or distinguishing attack [1].

The attacker can use the public keys  $\mathbf{P}_a = \langle \lfloor \frac{p}{q} \langle \mathbf{A} \mathbf{S}_a \rangle_q \rfloor \rangle_p$  and  $\mathbf{P}_b = \langle \lfloor \frac{p}{q} \langle \mathbf{S}_b^\top \mathbf{A} \rangle_q \rfloor \rangle_p$  to obtain information on  $\mathbf{S}_a$  and  $\mathbf{S}_b$ , respectively. We work out how this is done for  $\mathbf{S}_a$ , the case for  $\mathbf{S}_b$  being similar. Let  $1 \leq i \leq n$  and  $1 \leq j \leq \bar{n}$ . If we denote the  $i$ -th row of  $\mathbf{A}$  by  $\mathbf{a}_i^\top$  and the  $j$ -th column of  $\mathbf{S}_a$  by  $\mathbf{s}_j$ , then

$$(\mathbf{P}_a)_{i,j} = \langle \lfloor \frac{p}{q} \langle (\mathbf{a}_i, \mathbf{s}_j) \rangle_q \rfloor \rangle_p.$$

By the definition of the rounding function  $\lfloor \cdot \rfloor$ , we have that

$$(\mathbf{P}_a)_{i,j} \equiv \frac{p}{q} \langle (\mathbf{a}_i, \mathbf{s}_j) \rangle_q + e_{i,j} \pmod{p} \text{ with } e_{i,j} \in (-1/2, 1/2].$$

As  $\langle (\mathbf{a}_i, \mathbf{s}_j) \rangle_q = (\mathbf{a}_i, \mathbf{s}_j) + \lambda q$  for some integer  $\lambda$ , we infer that

$$\frac{q}{p} (\mathbf{P}_a)_{i,j} \equiv (\mathbf{a}_i, \mathbf{s}_j) + \frac{q}{p} e_{i,j} \pmod{q}. \quad (10)$$

So we have  $n$  equations involving  $\mathbf{s}_j$ . Unlike conventional LWE, the errors  $\frac{q}{p} e_{i,j}$  reside in a bounded interval, namely  $(-\frac{q}{2p}, \frac{q}{2p}]$ . In what follows, we will only consider the case that  $p$  divides  $q$ .

**Primal Attack:** In (10), we write  $\mathbf{s}$  for  $\mathbf{s}_j$ , denote by  $\mathbf{b}$  the vector of length  $m$  with  $j$ -th component  $\frac{q}{p} (\mathbf{P}_a)_{i,j}$ , and  $\mathbf{A}_m$  for the matrix consisting of the  $m$  top rows of  $\mathbf{A}$ . We then have, for  $\mathbf{e} \in (-\frac{q}{2p}, \frac{q}{2p}]^m$

$$\mathbf{b} \equiv \mathbf{A}_m \mathbf{s} + \mathbf{e} \pmod{q} \quad (11)$$

so that  $\mathbf{v} = (\mathbf{s}^\top, \mathbf{e}^\top, 1)^\top$  is in the lattice  $\Lambda$  defined as

$$\Lambda = \{ \mathbf{x} \in \mathbb{Z}^{n+m+1} : (\mathbf{A}_m | \mathbf{I}_m | -\mathbf{b}) \mathbf{x} = \mathbf{0} \pmod{q} \} \quad (12)$$

of dimension  $d = n + m + 1$  and volume  $q^m$  [9, 2]. The attacker then searches for a short vector in  $\Lambda$  which hopefully equals  $\mathbf{v}$ , thus enabling him to recover the secret  $\mathbf{s}$ .

**Lattice Rescaling:** The vector  $\mathbf{v} = (\mathbf{s}^\top, \mathbf{e}^\top, 1)^\top \in \Lambda$  (see Eq. 12 for the lattice  $\Lambda$ 's definition) is unbalanced due to  $\|\mathbf{s}\| \ll \|\mathbf{e}\|$ . For exploiting this fact, a rescaling technique originally due to [4], and analyzed further in [14] and [1] is applied in order to balance the expected norms per component of  $\mathbf{s}$  and of  $\mathbf{e}$ . As  $\mathbf{s}$  is ternary and has  $h$  non-zero entries, the expected norm of each entry equals  $\frac{h}{n}$ . Approximating each entry of  $\mathbf{e}$  as being drawn uniformly from  $(-\frac{q}{2p}, \frac{q}{2p}]$ , its expected norm equals  $\frac{1}{12} \frac{q^2}{p^2}$ . We therefore choose the scaling factor  $\omega$  satisfying

$$\omega = \sigma' \cdot \sqrt{\frac{n}{h}}, \text{ where } \sigma'^2 = \frac{q^2}{12p^2} \quad (13)$$

We now consider the lattice

$$\Lambda_\omega = \{(\omega \cdot \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) \in (\omega \cdot \mathbb{Z})^n \times \mathbb{Z}^m \times \mathbb{Z} : (\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) \in \Lambda\} \quad (14)$$

of dimension  $n+m+1$  and volume  $q^m \omega^n$ , where the lattice  $\Lambda$  is defined in Eq. 12. The attacker then tries to find a short vector of the form  $\mathbf{v}_\omega = (\omega \cdot \mathbf{s}^\top, \mathbf{e}^\top, 1)^\top \in \Lambda_\omega$  that  $\Lambda_\omega$  contains. This is typically done by using a lattice reduction algorithm to obtain a reduced basis of  $\Lambda_\omega$ , the first vector of which will be the shortest of that basis due to a common heuristic.

If the quality of the lattice reduction is good enough, the reduced basis will contain  $\mathbf{v}_\omega$ . The attack success condition is as in [2] assuming that BKZ [12, 25] with block-size  $b$  is used as the lattice reduction algorithm. The vector  $\mathbf{v}_\omega$  will be detected if its projection  $\tilde{\mathbf{v}}_b$  onto the vector space of the last  $b$  Gram-Schmidt vectors of  $\Lambda$  is shorter than the  $(d-b)^{th}$  Gram-Schmidt vector  $\tilde{\mathbf{b}}_{d-b}$ , where  $d$  is the dimension of  $\Lambda$  [2, Sec. 6.3],[9]. The condition that must be satisfied for the primal attack to succeed is therefore:

$$\begin{aligned} \|\tilde{\mathbf{v}}_b\| &< \|\tilde{\mathbf{b}}_{d-b}\| \\ \text{i.e., } \|\tilde{\mathbf{v}}_b\| &< \delta^{2b-d-1} \cdot (\text{Vol}(\Lambda))^{\frac{1}{d}} \\ \text{where, } \delta &= ((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi\mathbf{e}})^{\frac{1}{2(b-1)}} \end{aligned} \quad (15)$$

The attack success condition (15) yields the following *security* condition that must be satisfied by the key-exchange parameters in order for SPKEX to remain secure against the primal attack:

$$\begin{aligned} \sqrt{(\omega^2 \cdot h + \sigma'^2 m) \cdot \frac{b}{n+m}} &\geq \delta^{2b-d-1} \cdot (q^{d-n-1} \omega^n)^{\frac{1}{d}} \\ \text{where, } \delta &= ((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi\mathbf{e}})^{\frac{1}{2(b-1)}}, \\ \omega &= \sigma' \cdot \sqrt{n/h}, \\ \sigma' &= (q/2\sqrt{3}p), \\ \text{and } d &= m+n+1. \end{aligned} \quad (16)$$

The running time of the BKZ lattice reduction is  $b \cdot 2^{cb}$  CPU clock cycles [9], where  $c$  is an experimental constant, and depends on the underlying sieving

algorithm used. In our later analysis, we will use three values for  $c$ , like in [9, Sec 4.1]. Classically, the best known constant is  $c_C = \log_2 \sqrt{3/2} \approx 0.292$ , which is provided by the sieving algorithm from [7]. In the quantum context, the best constant equals  $c_Q = \log_2 \sqrt{13/9} \approx 0.265$ , [20, Sec. 14.2.10]. Finally, we use the “paranoid” value  $c_P = \log_2 \sqrt{4/3} \approx 0.2075$ , see [9, Sec 4.1].

**Dual Attack:** The dual attack against LWE attempts to find a short vector  $(\mathbf{v}, \mathbf{w}) \in \mathbb{Z}^m \times \mathbb{Z}^n$  in the dual lattice

$$\Lambda^* = \{(\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^m \times \mathbb{Z}^n : \mathbf{A}\mathbf{x} = \mathbf{y} \pmod{q}\}. \quad (17)$$

This vector is used to construct the distinguisher  $z = \{(\mathbf{v}, \mathbf{b})\}_q$ . If  $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q}$ , then  $z = \{(\mathbf{v}, \mathbf{b})\}_q \equiv \{(\mathbf{w}, \mathbf{s}) + (\mathbf{v}, \mathbf{e})\}_q$ , and  $z$  is therefore small and distributed as a Gaussian.

For an LWR distribution with uniform rounding error  $\mathbf{e}'$  and corresponding variance  $\sigma'^2 = q^2/12p^2$ , the distinguisher checks whether  $z = \{(\mathbf{v}, \mathbf{b})\}_q$  is small. For a non-LWR distribution,  $z$  is distributed uniformly modulo  $q$ . For an LWR distribution,  $z$ 's distribution approaches a Gaussian distribution of zero mean and variance  $\|\mathbf{v}\|^2 \cdot \sigma'^2$  as the lengths of the vectors  $\mathbf{v}$  and  $\mathbf{e}'$  increase, due to the Central limit theorem. The maximal statistical distance between this Gaussian distribution and the uniform distribution modulo  $q$  is bounded by  $\epsilon \approx (1/\sqrt{2}) \exp(-2\pi^2(\|\mathbf{v}\| \cdot \sigma'/q)^2)$ , a more detailed derivation of this result can be found in Appendix B. The attacker uses the BKZ algorithm with block-size  $b$  that outputs a short vector of length  $\delta^{d-1} \cdot \text{Vol}(\Lambda^*)^{1/d}$ , where  $d = m + n$  is the dimension of the dual lattice  $\Lambda^*$ , and its volume is  $\text{Vol}(\Lambda^*) = q^n$ .

As the key is hashed, a small advantage  $\epsilon$  is not sufficient. As explained in [2], assuming BKZ with block size  $b$ , the attack must be repeated at least  $R = \max(1, 1/2^{0.2075b} \cdot \epsilon^2)$  times. The cost of using the dual attack to distinguish an LWR distribution from uniform, employing BKZ with block size  $b$  and root-Hermite factor  $\delta$ , using  $m$  samples, where the LWR distribution is generated from an LWR problem instance of dimension  $n$ , large modulus  $q$ , rounding modulus  $p$  is:

$$\begin{aligned} \text{Cost}_{\text{Dual attack}} &= (b \cdot 2^{cb}) \cdot \max(1, 1/(\epsilon^2 \cdot 2^{0.2075 \cdot b})), \\ \text{where,} \quad \epsilon &= \frac{1}{\sqrt{2}} \cdot \mathbf{e}^{-2\pi^2 \left(\frac{\|\mathbf{v}\| \cdot \sigma'}{q}\right)^2}, \\ \|\mathbf{v}\| &= \delta^{m+n-1} \cdot \left(\frac{q}{\sigma' \cdot \sqrt{n/h}}\right)^{1/(m+n)}, \\ \delta &= ((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi \mathbf{e}})^{\frac{1}{2(b-1)}} \\ \text{and} \quad \sigma' &= (q/2\sqrt{3}p). \end{aligned} \quad (18)$$

The first term in the cost above, i.e.,  $(b \cdot 2^{cb})$  is the cost of running BKZ lattice reduction with block-size  $b$ , where the constant  $c$  is the BKZ sieving exponent. Finally, note that the cost in Eq. 18 also accounts for the fact that the dual

attack can be adapted against the sparse-ternary LWR problem by rescaling the dual lattice using an appropriate scaling factor  $\omega$  [1] (similar to lattice rescaling in the primal attack), resulting in the following rescaled dual lattice:

$$\Lambda_\omega^* = \{(\mathbf{x}, \mathbf{y}/\omega) \in \mathbb{Z}^m \times \left(\frac{1}{\omega} \cdot \mathbb{Z}^n\right) : \mathbf{A}\mathbf{x} = \mathbf{y} \pmod{q}\} \quad (19)$$

This scales the volume of the dual lattice from  $q^n$  to  $(q/\omega)^n$ , which correspondingly scales the norm  $\|\mathbf{v}\|$  of the short vector  $\mathbf{v}$  and hence the distinguishing advantage.

## 5.2 Hybrid Attack

In this section, we consider a hybrid lattice reduction and meet-in-the-middle attack (henceforth called *hybrid attack*) originally due to [19] and targeting the NTRU [18] cryptosystem. We first describe the hybrid attack, using notation similar to that of [28], in a general form. Subsequently, we specialize the attack to spKEX. Finally, we describe a scaling approach to make the hybrid attack use the fact that the secret is small and sparse.

The hybrid attack aims to find the shortest vector  $\mathbf{v}$  in a lattice  $\Lambda'$  with a basis  $\mathbf{B}'$  of the form:

$$\mathbf{B}' = \begin{pmatrix} \mathbf{B} & \mathbf{C} \\ \mathbf{0} & \mathbf{I}_r \end{pmatrix} \in \mathbb{Z}^{d \times d} \quad (20)$$

where  $0 < r < d$  is the meet-in-the-middle dimension,  $\mathbf{B} \in \mathbb{Z}^{(d-r) \times (d-r)}$ , and  $\mathbf{C} \in \mathbb{Z}^{(d-r) \times r}$ . Let  $\mathbf{v} = (\mathbf{v}_l^\top \mathbf{v}_g^\top)^\top$  be a short vector in  $\Lambda'$ . The vector  $\mathbf{v}_g \in \mathbb{Z}^r$  will be recovered by guessing, while  $\mathbf{v}_l \in \mathbb{Z}^{d-r}$  will be recovered using lattice reduction techniques (solving BDD problems). As the columns of  $\mathbf{B}'$  generate  $\Lambda'$ , it follows, because of the special form of the basis  $\mathbf{B}'$ , that there exists a  $\mathbf{x} \in \mathbb{Z}^{d-r}$  such that

$$\mathbf{v} = \begin{pmatrix} \mathbf{v}_l \\ \mathbf{v}_g \end{pmatrix} = \mathbf{B}' \begin{pmatrix} \mathbf{x} \\ \mathbf{v}_g \end{pmatrix} = \begin{pmatrix} \mathbf{B} & \mathbf{C} \\ \mathbf{0} & \mathbf{I}_r \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{v}_g \end{pmatrix} = \begin{pmatrix} \mathbf{B}\mathbf{x} + \mathbf{C}\mathbf{v}_g \\ \mathbf{v}_g \end{pmatrix} \quad (21)$$

As  $\mathbf{v}_l$  is short,  $\mathbf{C}\mathbf{v}_g$  is close to  $-\mathbf{B}\mathbf{x}$ , a vector from the lattice  $\Lambda(\mathbf{B})$ . As explained in [28], the idea now is that if we can correctly guess  $\mathbf{v}_g$ , we can hope to find  $\mathbf{v}_l$  by using Babai's Nearest Plane algorithm [3]. This algorithm, given a basis  $\tilde{\mathbf{B}}$ , finds for every target vector  $\mathbf{t} \in \mathbb{R}^{d-r}$  a vector  $\mathbf{e} = \text{NP}_{\tilde{\mathbf{B}}}(\mathbf{t})$  such that  $\mathbf{t} - \mathbf{e} \in \Lambda(\tilde{\mathbf{B}})$ .

The cost for the hybrid attack thus is the sum of two components: the cost of finding a good basis  $\tilde{\mathbf{B}}$  for  $\Lambda(\mathbf{B})$ , and the cost of generating  $\text{NP}_{\tilde{\mathbf{B}}}(\mathbf{C}\mathbf{y})$  for all  $\mathbf{y}$  from a large set  $N$  of vectors of length  $r$  that (with high probability) contains  $\mathbf{v}_g$ . The latter cost may be reduced by using a Meet-in-the-middle approach [19] which reduces the number of calls to the Nearest Plane algorithm to the square root of the number of calls with a brute-force approach.

We analyze the hybrid attack similarly as in [18]. For each pair  $(r, m)$  with  $1 \leq r, m < n$ , we stipulate that the quality of the reduced basis  $\tilde{\mathbf{B}}$  is so high that

$\text{NP}_{\tilde{\mathbf{B}}}(\mathbf{v}_g) = \mathbf{v}_l$  with probability one. The condition, derived from [19, Lemma 1] is that the norm of the last Gram-Schmidt vector of  $\tilde{\mathbf{B}}$  is at least twice  $\|\mathbf{v}_l\|_\infty$ , see also [18]. We use the Geometric Series Assumption to approximate the norms of these vectors in terms of the Hermite constant  $\delta$ . The cost for obtaining a reduced basis with Hermite constant  $\delta$  is estimated as  $b2^{cb}$ , where  $b$  is such that  $\delta \left( (\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi e} \right)^{\frac{1}{2b-1}}$ , and for the sieving constant  $c$  we take the three values considered before. Moreover, we estimate the cost for the lattice decoding part to equal the number of invocations of the Nearest Plane Algorithm, which, following [18], we set to  $2^{\frac{1}{2}r \cdot H}$ , where  $H$  is the entropy of the distribution of each of the coordinates of the guessed vector  $\mathbf{v}_g$ . The number  $2^{r \cdot H}$  is an approximation to the number of typical vectors of length  $r$ ; the factor  $\frac{1}{2}$  is due to either the usage of the MTM technique, or the use of Grover's algorithm in the quantum case. Finally, we minimized the cost over all feasible pairs  $(r, m)$ .

The hybrid attack will be applied to the lattice

$$\Lambda' = \{\mathbf{x} \in \mathbb{Z}^{n+m+1} \mid (\mathbf{A}|\mathbf{I}_m| - \mathbf{b})\mathbf{x} \equiv 0 \pmod{q}\}.$$

We aim to find the short vector  $\mathbf{v} = (\mathbf{s}^T, \mathbf{e}^T, 1)^T$  in  $\Lambda'$ . In the hybrid attack, we guess  $\mathbf{v}_g$ , the  $r < n$  top entries of  $\mathbf{s}$ . We know that all entries of  $\mathbf{s}$ , and hence all entries of  $\mathbf{v}_g$ , are from  $\{-1, 0, 1\}$ , and that  $\mathbf{s}$  has  $h$  non-zero entries. Note that the  $(n-r)$  top entries of  $\mathbf{v}_l$  are the  $(n-r)$  bottom entries of  $\mathbf{s}$ . In order to exploit the smallness and sparseness of these entries, the  $n-r$  leftmost columns of the matrix  $\mathbf{B}$ , are multiplied with an appropriate scaling factor  $\omega$ . Calculated similar to § 5.1 by equalizing the *per-component* expected norms of the secret  $\mathbf{s}$  and LWR rounding error  $\mathbf{e}$ , we arrive at the same scaling factor  $\omega = \frac{q^2}{12p^2} \cdot \sqrt{\frac{n}{h}}$ . This then scales up the volume of the  $(n-r+m+1)$  dimensional lattice  $\Lambda$  generated by the basis  $\mathbf{B}$  by a factor  $\omega^{n-r}$ .

Our analysis of the hybrid attack allows us to obtain a rough estimate of its cost. Wunderer [28] gives an extensive runtime analysis of the hybrid attack. One of the aspects he takes into account is that the attacker chooses a larger value of  $\delta$ . This leads to a smaller cost (running time) for lattice reduction to obtain  $\tilde{\mathbf{B}}$ , but decreases the probability that  $\text{NP}_{\tilde{\mathbf{B}}}(\mathbf{v}_g) = \mathbf{v}_l$ , thereby increasing the expected cost (running time) of the part of the attack dealing with solving BDD problems. Also, he takes into account that the guesses for  $\mathbf{v}_g$  are generated such that the most likely candidates for  $\mathbf{v}_g$  occur early, thus reducing the expected number of calls to the nearest plane algorithm.

### 5.3 Choice of Hamming Weight

Albrecht *et al* [1] describe an attack towards LWE and LWR with sparse secrets  $\mathbf{s}$ . Since most rows of the public matrix  $\mathbf{A}$  become irrelevant while calculating the product  $\mathbf{A} \cdot \mathbf{s}$  for such secrets, Albrecht's attack ignores a random  $k \leq n$  number of components of  $\mathbf{s}$  and brings down the lattice dimension (and hence attack cost) during lattice reduction. Non-zero components in  $\mathbf{s}$  may be incorrectly ignored, and the attack must therefore be repeated  $1/(1 - \frac{h}{n})^k$  times so

that with high likelihood, at least one of the times indeed all  $k$  ignored components of  $\mathbf{s}$  indeed equal zero. [1] also proposes an extension of this attack that attempts to reduce the above inaccuracy, however this extension involves significantly more repetitions, hence we do not consider it. For the standard Albrecht attack, we estimate the cost (in bits) for a given Hamming weight  $h \leq n$ , as the number of repetitions each low cost attack is performed times the cost of the low-cost attack on a lattice of smaller dimension  $n - k$ :

$$\frac{\binom{n}{k}}{\binom{n-h}{k}} \times \text{Cost}_{\text{Lattice Reduction}}(n, k, h)$$

Here  $\text{Cost}_{\text{Lattice Reduction}}(n, k, h)$  is defined as

$$\min\{b \cdot 2^{cb} \mid b \in \mathbb{N}, \text{ there exists } m \in \mathbb{N} \text{ such that } m \leq n \text{ and (22) is satisfied.}\}$$

$$\sqrt{(\omega^2 \cdot h + \sigma'^2 m) \cdot \frac{b}{n+m}} < \delta^{2b-d-1} \cdot (q^{d-(n-k)-1} \omega^{n-k})^{\frac{1}{d}}$$

$$\text{where,} \quad \delta = ((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi e})^{\frac{1}{2(b-1)}}, \quad (22)$$

$$\omega = \sigma' \cdot \sqrt{(n-k)/h},$$

$$\sigma' = (q/2\sqrt{3}p),$$

and

$$d = m + n - k + 1.$$

The term  $b \cdot 2^{cb}$  represents the cost of running BKZ lattice reduction with block-size  $b$ . The attack runs on a LWE problem of dimension  $n - k$  with  $m \leq n$  samples. Condition 22, which is essentially Condition 15, ensures that such an attack has chances of succeeding. Note that although the above applies to the primal attack, a similar analysis is possible for the dual attack, in which case  $\text{Cost}_{\text{Lattice Reduction}}(n, k, h)$  is calculated as in Eq. 18, with the parameter  $n$  replaced by  $n - k$ .

This specialized attack only gives an advantage if an attacker is able to choose a  $k$  for which the total attack cost is less than a standard lattice-reduction attack on a lattice of dimension  $n$ . Figure 1 depicts our estimations of how this specialized attack (cost) scales with  $k$  for different Hamming weights, considering the primal attack and assuming the usage of BKZ. Our experiments considered values of  $n$  in the range  $510 \leq n \leq 910$ . The attack cost for sparse-ternary LWR is estimated to be equivalent to the cost of attacking ternary LWR, for a Hamming weight of  $h = 0.2n$ , beyond which this specialized attack's cost increases with  $h$ , quickly becoming unfeasible.

Furthermore, to ensure that an exhaustive, brute-force search of each secret-key vector in the secret-key using Grover's quantum search algorithm [17] has a cost of at least  $\lambda$  (in bits), the chosen Hamming weight should satisfy:

$$\sqrt{\binom{n}{h}} \cdot 2^h > 2^\lambda \quad (23)$$

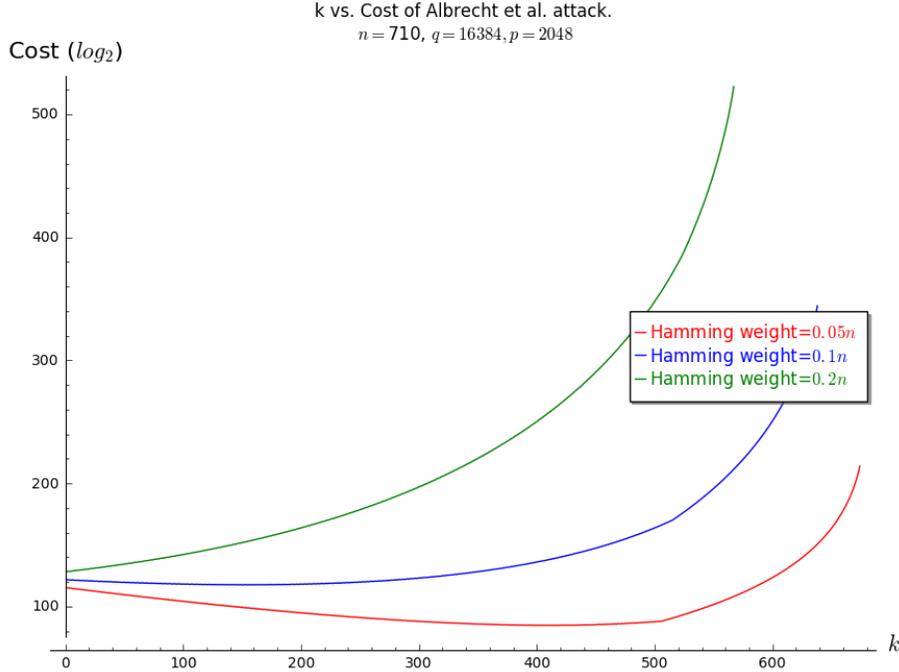


Figure 1: Cost of Albrecht *et al*'s [1] specialized (dual) attack against sparse-ternary LWR with Hamming weight  $0.05n$ ,  $0.1n$  and  $0.2n$ . Lowest cost for the two smaller Hamming weights is  $2^{85}$  and  $2^{118}$ , for  $k = 412$  and  $k = 151$  ignored secret components respectively. Lowest cost for Hamming weight  $0.2n$  is  $2^{128}$ , the same as non-sparse ternary LWR. The attack thus has no advantage against this Hamming weight.

Note that for a typical security level of  $\lambda = 128$ , a dimension of at least  $n = 512$  would be secure against Grover's quantum search, for any Hamming weight  $h$  that is at least  $0.1n$ .

#### 5.4 Unbiased Keys

spKEX's keys are unbiased. First,  $p$  is a divisor of  $q$  so that the rounding does not introduce any bias in the public-keys, as indicated in [6]. Moreover, as  $2^B$  divides  $p$ , there is no bias in  $\mathbf{K}$  either. As  $p$  is an integer multiple of  $2^{B+b_n+1}$ , the helper data does not leak information on  $\mathbf{K}$ .

#### 5.5 Choice and Computation of the Public Parameter $\mathbf{A}$

A fresh  $\mathbf{A}$  is computed in each protocol interaction (see Section 3) by permuting a fixed, public  $\mathbf{A}_{\text{master}}$  as  $\mathbf{A} = \Pi(\mathbf{A}_{\text{master}})$ . This is a trade-off between the two approaches from literature. Using a fixed  $\mathbf{A}$  (e.g., [22]) increases efficiency, but

opens up the possibility of pre-computation or backdoor-like attacks. Thus, [9] uses a fresh, random  $\mathbf{A}$  for *every* key-exchange session computed from a fresh random seed and a pseudo-random function. However, this causes a significant performance overhead [9], in particular, if a hardware accelerator for the pseudo-random function is not available. Such an approach is especially problematic in client-server scenarios where a single server engages in multiple key-exchange handshakes with multiple clients, and must generate  $\mathbf{A}$  for each client.

Security-wise, our approach ensures a sufficiently high randomness and freshness of  $\mathbf{A}$  preventing pre-computation attacks: a pre-computation attack based on lattice reduction is not feasible since the permutation would obtain a completely new lattice so that an attacker would be forced to perform lattice reduction on it. Performance-wise, our approach has a low overhead since permuting  $\mathbf{A}_{\text{master}}$  is computationally inexpensive and the random permutation can be efficiently encoded in a short random seed.

An example of a suitable permutation  $\Pi()$  randomly selects  $n_r \leq n$  rows in  $\mathbf{A}_{\text{master}}$  and shifts each of the selected rows by a random amount. This leads to a total of:

$$\binom{n}{n_r} \cdot n^{n_r} \quad (24)$$

possible permutations of  $\mathbf{A}_{\text{master}}$ . Considering a typical cryptographically large value of  $n$ , e.g., 512, and  $n_r = n$ , this leads to at least  $2^{256}$  possible random permutations  $\mathbf{A}$ .

## 5.6 Parameter Selection

Similar to [14] and [9], three sets of parameters are provided, namely classical, quantum and paranoid parameters. The parameters are chosen in such a way that the computational costs are at least  $2^{128}$  CPU cycles when running BKZ. The cost of running BKZ with block size  $b$  is estimated as  $b \cdot 2^{cb}$ , where for the classical, quantum and paranoid parameters we take  $c = c_C = 0.292$ ,  $c = c_Q = 0.265$  and  $c_P = 0.2075$ , respectively, see [9, Sec. 4.1]. For given parameters  $p, q, n$ , we minimize the attack cost over the block size  $b$  and the number of samples  $m$  that an attacker can use, such that the condition for successful attack is satisfied, for the primal, dual and specialized attack.

In all cases, we took the Hamming weight of the secret equal to  $0.2n$ . The chosen parameters are further optimized by scanning the values of  $b_h, p$ , and  $q$  that allow achieving the target security level for a failure probability lesser than  $2^{-40}$  and minimize bandwidth requirements. Finally, it is required that the length of the generated key, which is  $\bar{n} \cdot \bar{m} \cdot B$ , is at least equal to the desired key length, that is, at least 128 for the classical parameters, and at least 256 for the quantum and paranoid parameters. In order to avoid brute-force search over keys, we made sure that (23) is satisfied, with  $\lambda = 128$  for the classical parameters, and  $\lambda = 256$  for the quantum and paranoid parameters.

Table 2 shows the three sets of configuration parameters. We further compare them with related key-exchange and public-key encryption schemes based on the LWE and/or LWR problems, namely Frodo [9] and Lizard [14]. Note

Table 2: spKEX parameters and security level compared with Frodo [9] and Lizard [14]. Secrets have Hamming weight  $0.2n$ .

Case	$n$	$q$	$p$	$\sigma$	Primal (bits, b, m)	Dual (bits, b, m)	Hybrid (bits, r)
spKEX							
Classical	619 (LWR)	$2^{14}$	$2^9$	9.238 (LWR)	143, 460, 431	141, 451, 444	128, 278
Quantum	738 (LWR)	$2^{14}$	$2^{11}$	2.309 (LWR)	136, 480, 528	134, 472, 535	128, 278
Paranoid	864 (LWR)	$2^{14}$	$2^{11}$	2.309 (LWR)	130, 581, 621	128, 573, 596	135, 295
Frodo							
Classical	592 (LWE)	$2^{12}$	-	1.0 (LWE)	132, 442, 549	130, 438, 544	-
Quantum	752 (LWE)	$2^{15}$	-	1.323 (LWE)	132, 489, 716	130, 485, 737	-
Paranoid	864 (LWE)	$2^{15}$	-	1.323 (LWE)	129, 581, 793	128, 576, 833	-
Lizard							
Classical	544 (LWE), 840 (LWR)	$2^{10}$	$2^8$	5.988 (LWE), 1.155 (LWR)	?	128 (LWE), 132 (LWR)	?
Quantum	608 (LWE), 960 (LWR)	$2^{10}$	$2^8$	5.626 (LWE), 1.155 (LWR)	?	128 (LWE), 130 (LWR)	?
Paranoid	736 (LWE), 1450 (LWR)	$2^{10}$	$2^8$	6.4 (LWE), 1.155 (LWR)	?	128 (LWE), 129 (LWR)	?

Table 3: Proposed parameter sets for spKEX along with performance

Security level	Key size	$B, b_h$	Failure	$\log_2 q$	$n$	$\log_2 p$	$\sigma$	Bandwidth
Classical ( $c = 0.292$ )	128 bits	2, 3	$2^{-53}$	14	619	9	9.238	10.9 KB
Quantum ( $c = 0.265$ )	256 bits	4, 3	$2^{-42}$	14	738	11	2.309	15.88 KB
Paranoid ( $c = 0.2075$ )	256 bits	4, 4	$2^{-41}$	14	864	11	2.309	18.59 KB

that the authors of [9] calculate the security-levels of Frodo [9] considering that a discretized Gaussian error distribution is used, which might have a lower security than using an ideal Gaussian distribution. Lizard [14] does not disclose the optimal values of the BKZ block-size  $b$  and LWE or LWR samples  $m$  considered for calculating the security-levels mentioned in Table 2.

## 6 Performance Analysis

### 6.1 Bandwidth performance

Table 3 summarizes the proposed parameter sets introduced in Section 5.6. As mentioned before, we use a Hamming weight of  $0.2n$ , based on the reasoning in Section 5.3. The variance of noise introduced by LWR is computed as  $\sigma^2 = \frac{q^2}{12p^2}$ . The bandwidth requirements (in bits) are calculated as:

$$n \cdot (\bar{n} + \bar{m}) \cdot \lceil \log p \rceil + \bar{n} \cdot \bar{m} \cdot b_h \quad (25)$$

Table 4 compares spKEX with existing key-exchange, key-encapsulation and public-key encryption schemes based on the LWE and LWR problems. The use

Table 4: Comparison of SPKEX with other schemes for Quantum security parameters

Scheme	$n, q, p, \sigma$	Bandwidth	Failure
FRODO [9]	$752, 2^{15}, -, 1.323$	22.06 KB	$2^{-39}$
SPLWE-KEM [13]	$565, 520, -, 5.0$	158 KB	$2^{-7}$
LIZARD-PKE [14] (128-bit classical IND-CPA, 32-bit plaintext)	960 (LWR), 608 (LWE), $2^{10}, 2^8, 5.62$	32.88 KB	?
SPKEX	$738, 2^{14}, 2^{11}, 2.309$	15.88 KB	$2^{-42}$

of rounding has the highest impact on the public-key sizes and hence *reduces the bandwidth* requirements. The improved reconciliation method from [27] allows us a win-win-win scenario of *tolerating higher error* (as compared to Frodo), which results in (slightly) *higher security level*, and yet achieving a *lower failure probability* than Frodo (see Table 3). The improved error tolerance allowed us to further decrease  $p$ , thus dropping more bits from the public-keys, further *reducing their size and bandwidth* requirements. The usage of sparse ternary secrets leads to a higher value of  $n$  and thus increases bandwidth needs; however, as shown in the next section it has a positive bearing on the error, and on implementation.

## 6.2 Computational Performance

Three features of our key exchange improve the computational performance compared with related work.

*First*, the usage of a permutation to compute  $\mathbf{A}$  involves lower computational requirements. In particular, it requires the computation of  $n \lceil \log n \rceil$  bits compared with  $n^2 \lceil \log q \rceil / 2$  bits when all elements in  $\mathbf{A}$  are computed randomly resulting in an improvement of up to  $\frac{n \lceil \log q \rceil}{2 \lceil \log n \rceil}$  times. To illustrate this, we implement SPKEX using the row-wise permutation detailed in Section 5.5 and two variants of it, one in which  $\mathbf{A}$  is static and one in which  $\mathbf{A}$  is generated from a seed by means of a SHA3-based pseudo-random function. The row-wise permutation is computed from a 256-bit random seed by applying the same SHA3-based pseudo-random function on this seed to compute the shifts of each of the rows. The total key exchange for SPKEX is then performed in 3.657 msec. (fixed  $\mathbf{A}$ ), 7.774 msec. (permuted  $\mathbf{A}$ ), and 77.571 msec. (randomly generated  $\mathbf{A}$ ) respectively. Note that this performance numbers are obtained for three implementations that do not take advantage of the usage of sparse ternary secrets.

Table 5 details the performance of SPKEX on a 64-bit Intel Core i5-4210U CPU @ 1.70 GHz. Performance numbers illustrate the performance of SPKEX with the remaining design features: small rounding modulus  $p$  and sparse ternary secrets. The usage of small rounding modulus implies smaller operations and avoids the sampling of Gaussian noise, this speeds up both the computation of the public-keys and the final shared secret. The usage of sparse ternary

secrets has implementation advantages, simplifying multiplications of matrix components.

Table 5: CPU performance of SPKEX for quantum-security level

Key exchange phase	Time (msec.)
Alice’s secret-key	0.228
Alice public-key	1.211
Bob’s secret-key	0.228
Bob’s public-key	1.065
Bob’s shared key	0.014
Alice’s shared key	0.014
Total handshake	2.762

## 7 Conclusions & Future Work

We have presented SPKEX, a key-exchange scheme based on the sparse ternary secrets and Learning with Rounding. This combination seems to provide the best performance when bandwidth and implementation is considered as a whole. We further use the improved reconciliation method in [27] to achieve smaller public-key sizes, smaller bandwidth requirements, tolerance of higher error and a better failure probability. Finally, our approach to refresh the public matrix  $\mathbf{A}$  improves computational performance of existing approaches [9, 2] while preventing pre-computation attacks.

The usage of all these optimizations in spKEX together allows reducing the bandwidth requirements of Frodo around 30% while having significant implementation advantages.

An interesting direction of future work is the design of other cryptographic primitives such as key-encapsulation (with ephemeral secrets) based on similar design principles. Further research directions are the security analysis of the scheme and optimized implementations resistant to side-channel attacks.

## 8 Acknowledgments

We thank Zhenfei Zhang for fruitful discussions on the hybrid attack.

## References

- [1] Martin R. Albrecht. On dual lattice attacks against small-secret LWE and parameter choices in HELib and SEAL. Cryptology ePrint Archive, Report 2017/047, 2017. <http://eprint.iacr.org/2017/047>.

- [2] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - a new hope. *Cryptology ePrint Archive*, Report 2015/1092, 2015. <http://eprint.iacr.org/2015/1092>.
- [3] L. Babai. On lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
- [4] Shi Bai and Steven D. Galbraith. Lattice Decoding Attacks on Binary LWE. *Cryptology ePrint Archive*, Report 2013/839, 2013. <http://eprint.iacr.org/2013/839>.
- [5] Shi Bai, Adeline Langlois, Tancrede Lepoint, Amin Sakzad, Damien Stehle, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: using the Rényi divergence rather than the statistical distance. *Cryptology ePrint Archive*, Report 2015/483, 2015. <http://eprint.iacr.org/2015/483>.
- [6] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. *Cryptology ePrint Archive*, Report 2011/401, 2011. <http://eprint.iacr.org/2011/401>.
- [7] Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. *Cryptology ePrint Archive*, Report 2015/1128, 2015. <http://eprint.iacr.org/2015/1128>.
- [8] Andrej Bogdanov, Siyao Guo, Daniel Masny, Silas Richelson, and Alon Rosen. On the hardness of learning with rounding over small modulus. *Cryptology ePrint Archive*, Report 2015/769, 2015. <http://eprint.iacr.org/2015/769>.
- [9] Joppe Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! Practical, Quantum-Secure Key Exchange from LWE. *Cryptology ePrint Archive*, Report 2016/659, 2016. <http://eprint.iacr.org/2016/659>.
- [10] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 575–584, New York, NY, USA, 2013. ACM.
- [11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient Fully Homomorphic Encryption from (Standard) LWE. *Cryptology ePrint Archive*, Report 2011/344, 2011. <http://eprint.iacr.org/2011/344>.
- [12] Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better Lattice Security Estimates. In *Proceedings of the 17th International Conference on The Theory and Application of Cryptology and Information Security*, ASIACRYPT'11, pages 1–20, Berlin, Heidelberg, 2011. Springer-Verlag.

- [13] Jung Hee Cheon, Kyoo Hyung Han, Jinsu Kim, Changmin Lee, and Yongha Son. A Practical Post-Quantum Public-Key Cryptosystem Based on spLWE. Cryptology ePrint Archive, Report 2016/1055, 2016. <http://eprint.iacr.org/2016/1055>.
- [14] Jung Hee Cheon, Duhyeong Kim, Joohee Lee, and Yongsoo Song. Lizard: Cut off the Tail! Practical Post-Quantum Public-Key Encryption from LWE and LWR. Cryptology ePrint Archive, Report 2016/1126, 2016. <http://eprint.iacr.org/2016/1126>.
- [15] W. Diffie and M. Hellman. New directions in cryptography. *IEEE Trans. Inf. Theor.*, 22(6):644–654, September 1976.
- [16] Jintai Ding, Xiang Xie, and Xiaodong Lin. A simple provably secure key exchange scheme based on the learning with errors problem. Cryptology ePrint Archive, Report 2012/688, 2012. <http://eprint.iacr.org/2012/688>.
- [17] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 212–219, New York, NY, USA, 1996. ACM.
- [18] Jeff Hoffstein, Jill Pipher, John M. Schanck, Joseph H. Silverman, William Whyte, and Zhenfei Zhang. Choosing Parameters for NTRUEncrypt. Cryptology ePrint Archive, Report 2015/708, 2015. <http://eprint.iacr.org/2015/708>.
- [19] Nick Howgrave-Graham. *A Hybrid Lattice-Reduction and Meet-in-the-Middle Attack Against NTRU*, pages 150–169. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007.
- [20] Thijs Laarhoven. *Search problems in cryptography*. PhD thesis, Eindhoven University of Technology, 2015.
- [21] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: Extended abstract. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 333–342, New York, NY, USA, 2009. ACM.
- [22] Chris Peikert. Lattice cryptography for the internet. Cryptology ePrint Archive, Report 2014/070, 2014. <http://eprint.iacr.org/2014/070>.
- [23] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 84–93, New York, NY, USA, 2005. ACM.
- [24] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, February 1978.

- [25] C. P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Math. Program.*, 66(2):181–199, September 1994.
- [26] Peter W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.*, 26(5):1484–1509, October 1997.
- [27] Ludo Tolhuizen, Ronald Rietman, and Oscar Garcia-Morchon. Improved key-reconciliation method. Cryptology ePrint Archive, Report 2017/295, 2017. <http://eprint.iacr.org/2017/295>.
- [28] Thomas Wunderer. Revisiting the hybrid attack: Improved analysis and refined security estimates. Cryptology ePrint Archive, Report 2016/733, 2016. <http://eprint.iacr.org/2016/733>.

## A Upper bound on Failure Probability

In this section, we present an upper bound on the probability that Alice and Bob do not agree on a common key. In order to do so, we derive an upper bound on the probability that for given  $i, j$ , entry  $(i, j)$  of  $\{\mathbf{K}_a - \mathbf{K}_b\}_p$  is at least  $\Delta + 1$ . Note that  $\mathbf{K}_a - \mathbf{K}_b \equiv \mathbf{E}_b \cdot \mathbf{S}_a - \mathbf{S}_b^\top \cdot \mathbf{E}_a \pmod{p}$ . We consider entry  $(i, j)$  of  $\mathbf{K}_a - \mathbf{K}_b$ . We denote the  $i$ -th row of  $\mathbf{E}_b$  by  $\mathbf{e}'$ , the  $j$ -th column of  $\mathbf{S}_a$  by  $\mathbf{s}$ , the  $i$ -th row of  $\mathbf{S}_b^\top$  by  $\mathbf{s}'$ , and the  $j$ -th column of  $\mathbf{E}_a$  by  $\mathbf{e}$ . Note that  $\mathbf{s}$  and  $\mathbf{s}'$  are ternary vectors of length  $n$  and both have Hamming weight  $h$ . We are thus interested in finding an upper bound for

$$\Pr[|(\mathbf{s}, \mathbf{e}') - (\mathbf{s}', \mathbf{e})| \geq \Delta + 1.] \tag{26}$$

From the definition, it follows that  $\mathbf{e}'$  is determined by  $\mathbf{A}$  and  $\mathbf{s}'$ , and that  $\mathbf{e}$  is determined by  $\mathbf{A}$  and  $\mathbf{s}$ . We can assume that  $\mathbf{e}'$  and  $\mathbf{s}$  are independent, as well as  $\mathbf{e}$  and  $\mathbf{s}'$ . The inner products  $(\mathbf{s}, \mathbf{e}')$  and  $(\mathbf{s}', \mathbf{e})$ , however, are *not* independent. Indeed,  $\mathbf{e}$  has non-integer coefficients, so that  $(\mathbf{s}', \mathbf{e})$  need not be integer; similarly the inner product  $(\mathbf{s}, \mathbf{e}')$  need not be integer. The difference between these two inner products however, is integer. We take a conservative approach to work around this dependence.

By the triangle inequality,  $|(\mathbf{s}, \mathbf{e}') - (\mathbf{s}', \mathbf{e})|$  can only be at least  $\Delta + 1$  if at least one of  $|(\mathbf{s}, \mathbf{e}')|$  and  $|(\mathbf{s}', \mathbf{e})|$  is at least  $(\Delta + 1)/2$ . It follows that

$$\begin{aligned} \Pr[|(\mathbf{s}, \mathbf{e}') - (\mathbf{s}', \mathbf{e})| \geq \Delta + 1] &\leq \Pr[|(\mathbf{s}, \mathbf{e}')| \geq \frac{1}{2}(\Delta + 1) \vee |(\mathbf{s}', \mathbf{e})| \geq \frac{1}{2}(\Delta + 1)] \\ &\leq \Pr[|(\mathbf{s}, \mathbf{e}')| \geq \frac{1}{2}(\Delta + 1)] + \Pr[|(\mathbf{s}', \mathbf{e})| \geq \frac{1}{2}(\Delta + 1)]. \end{aligned}$$

We now assume that the entries of  $\mathbf{e}'$  are independently and uniformly distributed on  $(-1/2, 1/2]$ ; note that  $\mathbf{s}$  is not used in the computation of  $\mathbf{e}'$ . Similarly, we assume that the entries of  $\mathbf{e}$  are independently and uniformly

distributed on  $(-1/2, 1/2]$ . From the symmetry of the distribution of  $e'$  around 0, and the fact that  $\mathbf{s}$  has entries from  $\{-1, 0, 1\}$ , it follows that  $(\mathbf{s}, e')$  is distributed as the sum of  $h$  uniform variables on  $(-1/2, 1/2]$ , where  $h$  is the number of non-zero entries of  $\mathbf{s}$ . We thus wish to find the upper bound for the probability that a sum of  $h$  uniform variables on  $(-1/2, 1/2]$  is large. To do so, we apply the Chernoff bound.

If  $X_1, \dots, X_m$  are independent variables, then according to the Chernoff bound, for each  $a \geq 0$

$$\Pr\left(\sum_{i=1}^m X_i \geq a\right) \leq \min_{t>0} e^{-ta} \prod_{i=1}^m E[e^{tX_i}]. \quad (27)$$

If each  $X_i$  is uniformly distributed on  $(-1/2, 1/2]$ , then for each  $t > 0$

$$E[e^{tX_i}] = \int_{-1/2}^{1/2} e^{tx} dx = \frac{1}{t}(e^{t/2} - e^{-t/2}) \quad (28)$$

For each  $t > 0$ , we thus have that

$$\Pr\left(\sum_{i=1}^m X_i \geq a\right) \leq e^{-ta} \left(\frac{e^{t/2} - e^{-t/2}}{t}\right)^m. \quad (29)$$

As each  $X_i$  is in  $(-1/2, 1/2]$ , the only interesting case is  $a \leq \frac{1}{2}m$ . In (29), we write  $a = \frac{1}{2}\alpha m$ , set  $\tau = t/2$ , and obtain that for each  $\tau > 0$

$$\frac{1}{m} \ln\left(\Pr\left(\sum_{i=1}^m X_i \geq \frac{1}{2}\alpha m\right)\right) \leq -\alpha\tau + \ln\left(\frac{e^\tau - e^{-\tau}}{2\tau}\right). \quad (30)$$

Hence, if we define

$$f(\alpha) := \frac{1}{\ln 2} \min \left\{ -\alpha\tau + \ln\left(\frac{e^\tau - e^{-\tau}}{2\tau}\right) \mid \tau > 0 \right\},$$

then for each  $a > 0$

$$\Pr\left[\sum_{i=1}^m X_i > a\right] \leq 2^{m \cdot f\left(\frac{2a}{m}\right)}.$$

By differentiating to  $\tau$ , we infer that for fixed  $\alpha > 0$ , the function  $-\alpha\tau + \ln\left(\frac{e^\tau - e^{-\tau}}{2\tau}\right)$  is minimized for the solution  $\tau^*(\alpha)$  to the equation

$$\frac{e^\tau + e^{-\tau}}{e^\tau - e^{-\tau}} - \frac{1}{\tau} = \alpha.$$

We have not been able to find an analytical expression for  $\tau^*(\alpha)$  nor for  $f(\alpha)$ . The minimization can be carried out easily in SAGE by running `find_local_minimum((ln(sinh(x)/x)-(alpha x))/ln(2.0), 0, 5)`. Here (0,5) denotes the interval over which to minimize. In case the location of the maximum

is found at  $x = 5$ , it is best to increase the end point of the interval, and see if the location of the maximum changes.

Because of the symmetry of the probability distribution of  $X_1, \dots, X_m$ , we have that  $\Pr(|\sum_{i=1}^m X_i| \geq a) = 2(\Pr \sum_{i=1}^m X_i > a)$ . Combining all the above, with  $m = h$  and  $a = \frac{1}{2}(\Delta + 1)$ , we infer that the probability that Alice and Bob do not agree on entry  $(i, j)$  of the key is at most

$$4 \cdot 2^{h \cdot f((\Delta+1)/h)}.$$

The union bound implies that the probability that Alice and Bob do not agree on some entry of the key is at most  $\bar{m} \cdot \bar{n} \cdot (4 \cdot 2^{h \cdot f((\Delta+1)/h)})$ .

## B Distinguishing Advantage in Dual Attack

The dual attack against the Learning with Rounding (LWR) problem aims at distinguishing the LWR distribution from a uniform distribution modulo  $q$ . The advantage of the attack, i.e., the distinguishing advantage can be directly expressed in terms of the statistical distance between these two distributions. In general, for two probability distributions  $p_0$  and  $p_1$  on some domain  $D \subset \mathbb{R}$ , a distinguishing game can be described in terms of a challenger who chooses a uniformly random bit  $b \in \{0, 1\}$ , draws a random number  $x \in D$  according to distribution  $p_b$ , and reveals  $x$  to the adversary, who must then guess  $b$ . The advantage  $\epsilon$  of the adversary in winning this game can be expressed in terms of the advantage in distinguishing between the probability distributions  $p_0$  and  $p_1$ .

$$\epsilon = \frac{1}{2} \int_D |p_0(x) - p_1(x)| dx$$

For the case that  $p_0$  is a Gaussian distribution:  $p_0 = G_q(x, s^2)$  and  $p_1$  is the uniform distribution on  $(-q/2, q/2]$ , i.e.,  $p_1 = 1/q$  the advantage is

$$\begin{aligned} \epsilon &= \frac{1}{2} \int_{-q/2}^{q/2} \left| \sum_{k=-\infty}^{\infty} \frac{1}{\sqrt{2\pi s^2}} e^{-(x-kq)^2/2s^2} - \frac{1}{q} \right| dx \\ &= \frac{1}{2} \int_{-1/2}^{1/2} \left| \sum_{k=-\infty}^{\infty} \frac{\alpha}{\sqrt{2\pi}} e^{-\alpha^2(\xi-k)^2/2} - 1 \right| d\xi, \quad \text{where } \alpha = q/s. \end{aligned}$$

$$\epsilon \approx \int_{-a}^a \frac{\alpha}{\sqrt{2\pi}} e^{-\alpha^2 \xi^2/2} - 1 \, d\xi = \operatorname{erf}\left(\frac{\alpha a}{\sqrt{2}}\right) - 2a = \operatorname{erf}\left(\ln\left(\frac{\alpha}{\sqrt{2\pi}}\right)\right) - \frac{2}{\alpha} \sqrt{\ln\left(\frac{\alpha^2}{2\pi}\right)}$$

For the general case however, i.e., as the standard deviation  $s$  of the Gaussian distribution increases,  $\alpha$  decreases and the terms where  $k \neq 0$  must also be considered in the summation. We can find an approximate upper bound for the

advantage using the Cauchy-Schwarz inequality as:

$$\begin{aligned}
\epsilon &= \frac{1}{2} \int_{-1/2}^{1/2} \left| \sum_{k=-\infty}^{\infty} \frac{\alpha}{\sqrt{2\pi}} e^{-\alpha^2(\xi-k)^2/2} - 1 \right| d\xi \\
&\leq \frac{1}{2} \left( \int_{-1/2}^{1/2} \left( \sum_{k=-\infty}^{\infty} \frac{\alpha}{\sqrt{2\pi}} e^{-\alpha^2(\xi-k)^2/2} - 1 \right)^2 d\xi \right)^{1/2} \\
&= \frac{1}{2} \left( \int_{-1/2}^{1/2} \left( \sum_{k=-\infty}^{\infty} \frac{\alpha}{\sqrt{2\pi}} e^{-\alpha^2(\xi-k)^2/2} \right)^2 d\xi - 1 \right)^{1/2} \\
&= \frac{1}{2} \left( \sum_{m=-\infty}^{\infty} |\hat{\phi}_m|^2 - 1 \right)^{1/2} \quad (\text{using Parseval's theorem}).
\end{aligned}$$

The numbers  $\hat{\phi}_m$  are the Fourier components of the periodic function  $\xi \rightarrow qG_s(q\xi, s^2)$ :

$$\hat{\phi}_m = \int_{-1/2}^{1/2} e^{-2\pi im\xi} \sum_{k=-\infty}^{\infty} \frac{\alpha}{\sqrt{2\pi}} e^{-\alpha^2(\xi-k)^2/2} d\xi = e^{-2\pi^2 m^2/\alpha^2}$$

This gives the following upper bound for the distinguishing advantage  $\epsilon$ :

$$\epsilon \leq \frac{1}{\sqrt{2}} \left( \sum_{m=1}^{\infty} e^{-4\pi^2 m^2/\alpha^2} \right)^{1/2}$$

For general standard deviation  $s$  and large  $\alpha$ , the  $m = 1$  term dominates the summation, giving the approximate upper bound of  $\epsilon \leq \exp(2\pi^2 s^2/q^2)/\sqrt{2}$ .