Classical Reduction of SVP to LWE: A Concrete Security Analysis

Palash Sarkar and Subhadip Singha
Applied Statistics Unit
Indian Statistical Institute
203, B.T.Road, Kolkata, India - 700108.
{palash, subha_r}@isical.ac.in

July 12, 2020

Abstract

Regev (2005) introduced the learning with errors (LWE) problem and showed a quantum reduction from a worst case lattice problem to LWE. Building on the work of Peikert (2009), a classical reduction from the shortest vector problem to LWE was obtained by Brakerski et al. (2013). A concrete security analysis of Regev's reduction by Chatterjee et al. (2016) identified a huge tightness gap. The present work performs a concrete analysis of the tightness gap in the classical reduction of Brakerski et al. It turns out that the tightness gap in the Brakerski et al. classical reduction is even larger than the tightness gap in the quantum reduction of Regev. This casts doubts on the implication of the reduction to security assurance of practical cryptosystems.

Keywords: lattices, shortest vector problem, learning with errors, classical reduction, concrete analysis.

Mathematics Subject Classification: Primary: 94A60

1 Introduction

In a landmark paper, Regev [14] introduced the learning with errors (LWE) problem. Many cryptosystems have based their security on the hardness of the LWE problem. Examples of such cyptosystems are Frodo [2], Kyber [3], LAC [11], NewHope [1], Round5 [4] and Saber [8] all of which are candidates for standardisation as a post-quantum cryptosystem to be selected by the NIST of the USA. One reason for confidence in the hardness of the LWE problem is a reduction proved by Regev [14] from a worst-case lattice problem to LWE. The reduction obtained by Regev was quantum, i.e., the algorithm is required to make a quantum computation.

A problem left open by Regev was whether there is a classical reduction from a worst case lattice problem to LWE. The initial answer to this problem was provided by Peikert [13]. While this represented progress, Peikert's reduction was not considered to be satisfactory since either an exponential size modulus is required or, the lattice problem considered is not one of the standard problems. Later work by Brakerski et al. [6] built on Peikert's work to show a classical reduction from a standard lattice problem to LWE avoiding the exponential size modulus.

The works of Regev [14], Peikert [13] and Brakerski et al. [6] are all in the asymptotic setting where the lattice dimension is allowed to go to infinity. Practical cryptosystems, on the other hand, have a fixed value of the lattice dimension. So, it is of interest to know what kind of security assurance one obtains from the results of [14, 13, 6] for practical cryptosystems. Suppose it is believed that a lattice problem $\mathcal{P}$ is computationally hard. It is desired to translate this into a belief that a particular cryptosystem $\mathcal{C}$ is difficult to break, i.e., the difficulty of solving $\mathcal{P}$ is reduced to the difficulty of breaking $\mathcal{C}$. In other words, it is required to show that if there is an algorithm $\mathcal{A}$ to break $\mathcal{C}$, then there is an algorithm $\mathcal{B}$ (which uses $\mathcal{A}$ as an oracle) to solve $\mathcal{P}$. Suppose

$\mathcal{A}$ takes time T and has success probability P_S and further, $\mathcal{B}$ takes time T' and has success probability P'_S. The tightness gap of the reduction is defined to be $(T'/P'_S)/(T/P_S)$. The reduction is said to be tight if the tightness gap is one (or, small). On the other hand, if the tightness gap is very large, then the usefulness of the reduction for obtaining security assurance of a practical cryptosystem becomes questionable.

The tightness gap of the reduction given by Regev was first investigated in [7] and in more details in [16]. The results of [7, 16] indicate that the tightness gap is very large. Based on the analysis in [7], Bernstein [5] comments that "the loss of tightness is gigantic" in [14].

In this paper, we follow up on [7, 16] and perform a concrete security analysis of the tightness gap of the reduction in [6]. The reduction of Peikert [13] is a step in the reduction performed by Brakerski et al. [6]. As a first step, we work out the tightness gap of Peikert's reduction. Then we follow the proof strategy in Brakerski et al. [6] and finally work out the end-to-end tightness gap of the classical reduction from the shortest vector problem to the LWE. There are two aspects to the concrete analysis. The first is a quadratic loss in the dimension of the lattice and the second is a loss of tightness. The loss of tightness in this classical reduction is more than that of the original quantum reduction by Regev [14]. The quadratic loss in the dimension was already pointed out in [6]. Due to this quadratic loss, Brakerski et al. put forward the open question of obtaining a reduction without such a loss mentioning that this would amount to a full de-quantization of Regev's reduction. The paper [6], however, does not consider the issue of the loss in tightness. Our analysis shows that due to this loss of tightness, the reduction is not very meaningful in practice, especially for determining the sizes of the parameters of a cryptosystem which would purpotedly enjoy the protection offered by the hardness of a well studied worst case lattice problem.

2 Preliminaries

Fix a positive integer n. Let $\mathbf{B}$ be an $n \times n$ matrix whose columns are n linearly independent vectors in $\mathbb{R}^n$. The lattice $L = L(\mathbf{B})$ generated by $\mathbf{B}$ is the set of all vectors $\mathbf{Ba}$ where $\mathbf{a} = (a_1, \ldots, a_n)^\intercal \in \mathbb{Z}^n$. The columns of $\mathbf{B}$ (or, more generally $\mathbf{B}$ itself) is called a basis of the lattice L. Let $\mathbf{b}_1, \ldots, \mathbf{b}_n$ denote the columns of $\mathbf{B}$. The Gram-Schmidt orthogonalisation (GSO) of $\mathbf{b}_1, \ldots, \mathbf{b}_n$ will be denoted as $\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_n$.

The length of a vector in L will be considered to be given by its Euclidean norm. For $i \in \{1, \ldots, n\}$, let $\lambda_i(L)$ be the least real number r such that L has i linearly independent vectors with the longest having length r. In particular, we will be interested in $\lambda_1(L)$, which is the smallest possible length of any non-zero lattice vector.

The dual of a lattice L is denoted as L^* and is defined to be the set of all vectors $\mathbf{y} \in \mathbb{R}^n$ such that $\langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}$ for all $\mathbf{x} \in L$. Given a basis $\mathbf{B}$ for L, the matrix $\mathbf{B}^* = (\mathbf{B}^{-1})^\intercal$ is a basis for L^* and is called the dual basis of $\mathbf{B}$.

Since $(\mathbb{Z}, +)$ is a subgroup of $(\mathbb{R}, +)$, the quotient group $\mathbb{R}/\mathbb{Z}$ is represented by the interval $\mathbb{T} = [0, 1)$ with addition modulo 1. The cyclic subgroup $\{0, 1/p, \ldots, (p-1)/p\}$ of $\mathbb{T}$ of order p will be denoted by $\mathbb{T}_p$. The normal distribution with mean μ and standard deviation σ will be denoted as $\mathcal{N}(\mu, \sigma)$. For $\alpha \in (0, 1)$, Ψ_α is the probability distribution over $\mathbb{T}$ obtained by sampling from $\mathcal{N}(0, \alpha/\sqrt{2\pi})$ and reducing the result modulo 1.

Fix an integer $p \geq 2$. Let $\mathbf{s}$ be chosen uniformly at random from $\mathbb{Z}_p^n$. Let χ be a probability distribution on $\mathbb{Z}_p$. The distribution $A_{p,\mathbf{s},\chi}$ on $\mathbb{Z}_p^n \times \mathbb{Z}_p$ is defined as follows: choose $\mathbf{a}$ uniformly at random from $\mathbb{Z}_p^n$; e from $\mathbb{Z}_p$ following χ and output $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e)$, where the addition is performed modulo p. Let ϕ be a probability density function on $\mathbb{T}$. The distribution $A_{p,\mathbf{s},\phi}$ is defined as follows: choose $\mathbf{a}$ uniformly at random from $\mathbb{Z}_p^n$; e from $\mathbb{T}$ following ϕ and output $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle / p + e)$, where the addition is performed modulo 1. When $\phi = \Psi_\alpha$, the distribution $A_{p,\mathbf{s},\Psi_\alpha}$ is written more conveniently as $A_{p,\mathbf{s},\alpha}$.

For $\mathbf{x} \in \mathbb{R}^n$ and $s > 0$, define $\rho_s(\mathbf{x}) = \exp\left(-\pi ||\mathbf{x}||^2 / s^2\right)$. For a lattice L, define $\rho_s(L) = \sum_{\mathbf{x} \in L} \rho_s(\mathbf{x})$. The discrete Gaussian distribution $D_{L,s}$ on a lattice L assigns to a vector $\mathbf{v} \in L$ the probability $D_{L,s}(\mathbf{v}) = \rho_s(\mathbf{v})/\rho_s(L)$. For a lattice L and a real number $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(L)$ is the smallest s such that $\rho_{1/s}(L^* \setminus \{0\}) \leq \epsilon$.

The origin centered parallelopiped $\mathcal{P}_{1/2}(\mathbf{B})$ of a basis $\mathbf{B}$ is defined to be $\mathcal{P}_{1/2}(\mathbf{B}) = \{\mathbf{Bc} : \mathbf{c} \in [-1/2, 1/2)^n\}$.

For $\mathbf{w} \in \mathbb{R}^n$ and basis $\mathbf{B}$, the vector $\mathbf{x} = \mathbf{w} \bmod \mathbf{B}$ is the unique $\mathbf{x} \in \mathcal{P}_{1/2}(\mathbf{B})$ such that $\mathbf{w} - \mathbf{x} \in L(\mathbf{B})$; further, $\mathbf{x} = \mathbf{B}(\mathbf{B}^{-1}\mathbf{w} - \lfloor \mathbf{B}^{-1}\mathbf{w} \rfloor)$.

Let X be a random variable taking values in a set D and S be a subset of D. By $f_X(S)$ we denote the probability that X takes values in S. Given two random variables X and Y over D, the statistical distance between them is denoted as $\Delta(X, Y)$ and is defined to be $\Delta(X, Y) = \max_{S \subseteq D} |f_X(S) - f_Y(S)|$.

By $\mathcal{B}_n$ we will denote the open ball in $\mathbb{R}^n$ of unit radius, i.e., $\mathcal{B}_n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| < 1\}$. For a real number d and $\mathbf{z} \in \mathbb{R}^n$, the open ball in $\mathbb{R}^n$ centered at $\mathbf{z}$ and of radius d will be denoted as $\mathbf{z} + d \cdot \mathcal{B}_n$. The notation $\mathbf{w} \xleftarrow{\$} \mathbf{z} + d \cdot \mathcal{B}_n$ denotes the choice of a vector $\mathbf{w}$ drawn uniformly from $\mathbf{z} + d \cdot \mathcal{B}_n$.

2.1 Computational Problems

Let $\gamma(n) \geq 1$ be a function from the naturals to the naturals. The problem GapSVP_γ is the following: The input is a pair $(\mathbf{B}, d)$, where $\mathbf{B}$ is a basis of an n-dimensional lattice $L = L(\mathbf{B})$ and $d > 0$ is a real number. The input is a YES instance if $\lambda_1(L) \leq d$ and it is a NO instance if $\lambda_1(L) \geq \gamma(n) \cdot d$. The problem SIVP_γ is the following: The input is a basis $\mathbf{B}$ of an n-dimensional lattice $L = L(\mathbf{B})$ and the task is to output a set of n linearly independent vectors from L whose lengths are at most $\gamma(n) \cdot \lambda_n(L)$.

Let φ be a real valued function defined on lattices. The discrete Gaussian sampling (DGS_φ) problem is the following: The input is a pair $(\mathbf{B}, r)$, where $\mathbf{B}$ is a basis of an n-dimensional lattice $L = L(\mathbf{B})$ and $r > \varphi(L)$ is a real number. The task is to output a sample from $D_{L,r}$.

A variant of the closest vector problem (CVP) was considered in [15]: The input is $(\mathbf{B}, d, \mathbf{x})$, where $\mathbf{B}$ is the basis of an n-dimensional lattice $L = L(\mathbf{B})$, d is a positive real number with $d < \lambda_1(L)/2$, and $\mathbf{x} \in \mathbb{R}^n$. The task is to find the closest lattice point to $\mathbf{x}$ (since $d < \lambda_1(L)/2$, there is a unique closest vector). This problem is also the bounded distance decoding problem [13].

The learning with errors problem $\text{LWE}_{n,p,\chi}$ is the following. For uniform random $\mathbf{s}$ in $\mathbb{Z}_p^n$, given samples from $A_{p,s,\chi}$, it is required to output $\mathbf{s}$. If the number of samples is m, then the problem is denoted as $\text{LWE}_{n,m,p,\chi}$. Similarly, for a probability density function ϕ on $\mathbb{T}$, the $\text{LWE}_{n,m,p,\phi}$ problem is the following. For uniform random $\mathbf{s}$ in $\mathbb{Z}_p^n$, given samples from $A_{p,s,\phi}$, it is required to output $\mathbf{s}$. If the number of samples is m, then the problem is denoted as $\text{LWE}_{n,m,p,\phi}$. Both versions of the LWE problem were introduced by Regev in [15]. When $\phi = \Psi_\alpha$, the problem $\text{LWE}_{n,m,p,\phi}$ is more conveniently written as $\text{LWE}_{n,m,p,\alpha}$.

The problem ζ-to-γ-GapSVP (denoted as $\text{GapSVP}_{\zeta,\gamma}$) was introduced in [13]. For functions $\zeta(n) \geq \gamma(n) \geq 1$, an input to $\text{GapSVP}_{\zeta,\gamma}$ is a pair $(\mathbf{B}, d)$, where $\mathbf{B}$ is a basis of an n-dimensional lattice $L = L(\mathbf{B})$ for which $\lambda_1(L) \leq \zeta(n)$, $\min_i \|\widetilde{\mathbf{b}}_i\| \geq 1$, and $1 \leq d \leq \zeta(n)/\gamma(n)$. The input is a YES instance if $\lambda_1(L) \leq d$ and it is NO instance if $\lambda_1(L) > \gamma(n) \cdot d$. It has been shown in [13] that for $\zeta(n) \geq 2^{n/2}$, the $\text{GapSVP}_{\zeta,\gamma}$ problem is equivalent to the standard GapSVP_γ problem.

3 Reducing DGS to LWE

Regev [15] described a quantum algorithm which given access to an LWE oracle can solve the SIVP (or, the GapSVP). In the first step, the SIVP is reduced to the DGS problem using a classical algorithm. The main part of the proof is a quantum algorithm which reduces the DGS problem to the LWE problem. The proof given by Regev [15] is in an asymptotic setting. A concrete analysis of the tightness gap in the reduction was carried out in [7] and in more details in [16]. We provide a brief overview of Regev's DGS-to-LWE reduction using some of the terminology used in [16].

Let p be a positive integer and $\alpha \in (0, 1)$. Assume that an oracle $\mathsf{solveLWE}_{n,n^c,p,\Psi_\alpha}(\mathcal{I})$ is available for some constant $c > 0$. The input $\mathcal{I}$ to the oracle consists of n^c samples from $A_{p,\mathbf{s},\Psi_\beta}$ for some $0 < \beta \leq \alpha$. The oracle is guaranteed to work correctly if $\beta = \alpha$, otherwise it might return an incorrect result. Let $\mathbf{B}$ be an $n \times n$ basis matrix of an n-dimensional lattice $L = L(\mathbf{B})$ and r is a real number satisfying $r \geq \sqrt{2n} \cdot \eta_\epsilon(L)/\alpha$. The goal is

to design an algorithm $\mathsf{solveDGS}(\mathbf{B}, r)$ which returns a sample from $D_{L,r}$ using the oracle $\mathsf{solveLWE}_{n,n^c,p,\Psi_\alpha}(\mathcal{I})$ where $\alpha p > 2\sqrt{n}$.

Let $r_i = r \cdot (\alpha p/\sqrt{n})^i$ for $i = 1, \ldots, 3n$. A list $\mathcal{L}$ containing samples from $D_{L,r_{3n}}$ can be created without using the LWE oracle. The algorithm $\mathsf{solveDGS}(\mathbf{B}, r)$ starts with such a list and iterates a procedure over $3n$ steps with i going down from $3n$ to 1. The i-th step updates the list $\mathcal{L}$ consisting of n^c samples from D_{L,r_i} with n^c samples from $D_{L,r_{i-1}}$. At the end of the procedure, a sample from the final list $\mathcal{L}$ is returned. Each iteration updates the list $\mathcal{L}$ using a quantum sampling procedure n^c times. Each application of the quantum sampling procedure uses a classical algorithm $\mathsf{solveCVP}(L^*, \mathcal{L}, \mathbf{z})$, where L^* is the dual lattice of L, $\mathcal{L}$ contains n^c samples from D_{L,r_i} for some $i \in \{1, \ldots, 3n\}$, and $\mathbf{z}$ is within distance $\lambda_1(L^*)/2$ of L^*. The algorithm $\mathsf{solveCVP}$ solves the CVP problem for L^* mentioned in Section 2.1. It is the algorithm $\mathsf{solveCVP}$ which invokes the oracle $\mathsf{solveLWE}_{n,n^c,p,\Psi_\alpha}(\mathcal{I})$. So, the main part of the DGS-to-LWE reduction is the design of the algorithm $\mathsf{solveCVP}$.

In Regev's reduction, $\mathsf{solveCVP}(L^*, \mathcal{L}, \mathbf{z})$ solves the unique closest vector problem on L^* using a list $\mathcal{L}$ of samples from $D_{L,\mathfrak{r}}$ with $\mathfrak{r} \geq \sqrt{2}p \cdot \eta_\epsilon(L)$, and $\mathbf{z}$ is within distance $\alpha q/(\sqrt{2}r) < \lambda_1(L^*)/2$ of L^*. As used in [13], by interchanging the roles of L and L^*, it is possible to invoke $\mathsf{solveCVP}(L, \mathcal{L}, \mathbf{z})$ to solve the unique closest vector problem on L using a list $\mathcal{L}$ of samples from $D_{L^*,\mathfrak{r}}$ with $\mathfrak{r} \geq \sqrt{2}p \cdot \eta_\epsilon(L^*)$, and $\mathbf{z}$ is within distance $\alpha q/(\sqrt{2}r) < \lambda_1(L)/2$ of L. We record this as follows.

Proposition 1. *[15, 13] Let $\mathbf{B}$ be an $n \times n$ basis matrix for an n-dimensional lattice $L = L(\mathbf{B})$, p be a positive integer, $\mathfrak{r}$ be a real number satisfying $\mathfrak{r} \geq \sqrt{2}p \cdot \eta_\epsilon(L^*)$ and $\alpha \in (0, 1)$ be such that $\alpha p > 2\sqrt{n}$. Let $c > 0$ be a constant. Given a list $\mathcal{L}$ consisting of n^c samples from $D_{L^*,\mathfrak{r}}$ and an oracle $\mathsf{solveLWE}_{n,n^c,p,\Psi_\alpha}(\mathcal{I})$, where $\mathcal{I}$ consists of n^c samples from $A_{p,\mathbf{s},\Psi_\beta}$ for some $0 < \beta \leq \alpha$, there is an algorithm $\mathsf{solveCVP}(L, \mathcal{L}, \mathbf{z})$, where $\mathbf{z}$ is within distance $\alpha q/(\sqrt{2}r) < \lambda_1(L)/2$ of L, which finds the unique vector in L which is closest to $\mathbf{z}$.*

Following [16], we have the following facts.

1. Algorithm $\mathsf{solveCVP}$ calls the oracle $\mathsf{solveLWE}$ a total of n^{2c+2} times.
2. The success probability of algorithm $\mathsf{solveCVP}$ is at least

$$\left(1 - \max(\exp\left(-m(\mu_0 - t)^2/2\right), \exp\left(-mt^2/2\right))\right)^{n^{2c+2}} \tag{1}$$

where $\mu_0 = \exp(-\pi\alpha^2)$, and $t \in (0, \mu_0)$ and $m \leq n^c$ are chosen so as to maximise (1). Setting $m = n^c$ and $t = \mu_0/2$, the expression in (1) becomes

$$\left(1 - \exp\left(-n^c \exp\left(-2\pi\alpha^2\right)/8\right)\right)^{n^{2c+2}} \tag{2}$$

Using this lower bound for the success probability, it has been shown in [16] that an upper bound on the tightness gap of the DGS to LWE reduction is the following.

$$3n^{3c+3} \cdot \left(1 - \exp\left(-n^c \exp(-2\pi\alpha^2)/8\right)\right)^{-3n^{3c+3}}. \tag{3}$$

For most practical cryptosystems[1], α is at most $1/\sqrt{n}$. Considering $\alpha = 1/\sqrt{n}$, the tightness gap given by (3) is essentially $3n^{3c+3}$ [16]. The tightness gap of the reduction from DGS to LWE has been extended to obtain the tightness gap of the reduction from SIVP to average-case decision LWE in [7] and updated in [16] and is given by the following expression.

$$6pn^{3c+d_1+2d_2+9}. \tag{4}$$

[1] This was mentioned by Chris Peikert in an email.

4 Reducing GapSVP$_{\zeta,\gamma}$ to LWE

Peikert [13] showed a classical reduction of GapSVP$_{\zeta,\gamma}$ to LWE$_{n,n^c,q,\Psi_\alpha}$, where $\gamma = \gamma(n) \geq n/(\alpha\sqrt{\log n})$, $q = q(n) \geq \zeta(n) \cdot \omega(\sqrt{\log n/n})$ and $c > 0$ is a constant. The reduction makes use of Proposition 1, i.e., it uses an LWE oracle to solve CVP.

Let $\mathbf{B}$ be an $n \times n$ basis matrix of an n-dimensional lattice $L = L(\mathbf{B})$ and $r \geq \max_i \|\widetilde{b}_i\| \cdot \omega(\sqrt{\log n})$. By $\mathsf{sample}(\mathbf{B}, r)$ we denote the sampling algorithm which on input $\mathbf{B}$ and r returns a sample which is within negligible statistical distance from $D_{L,r}$. Such an algorithm is described in [9].

The algorithm for reducing GapSVP$_{\zeta,\gamma}$ to LWE given by Peikert [13] is shown in Algorithm 1. The algorithm $\mathsf{solveCVP}$ in turn calls the LWE oracle $\mathsf{solveLWE}$. So, overall $\mathsf{solveGapSVP}_{\zeta,\gamma}$ solves GapSVP$_{\zeta,\gamma}$ by calling the LWE oracle $\mathsf{solveLWE}$. Algorithm $\mathsf{solveGapSVP}_{\zeta,\gamma}$ calls $\mathsf{solveCVP}$ a total of N times.

Algorithm 1 Reducing GapSVP$_{\zeta,\gamma}$ to LWE$_{q,\Psi_\alpha}$, where $\gamma = \gamma(n) \geq n/(\alpha\sqrt{\log n})$ and $q = q(n) \geq \zeta(n) \cdot \omega(\sqrt{\log n/n})$.

function $\mathsf{solveGapSVP}_{\zeta,\gamma}(\mathbf{B}, d)$
  Let $\mathbf{D}$ be the reverse dual basis of $\mathbf{B}$;
  $d' = d \cdot \sqrt{n/(4 \ln n)}$; $r = q\sqrt{2n}/(\gamma d)$;
  for $i \leftarrow 1$ to N **do**
    $\mathbf{w} \xleftarrow{\$} d' \cdot \mathcal{B}_n$; $\mathbf{x} = \mathbf{w} \bmod \mathbf{B}$;
    $\mathcal{L} \leftarrow \{\}$;
    for $j \leftarrow 1$ to n^c **do**
      $\mathcal{L} \leftarrow \mathcal{L} \cup \mathsf{sample}(D, r)$;
    end for
    $\mathbf{v} \leftarrow \mathsf{solveCVP}(\mathbf{B}, \mathcal{L}, \mathbf{x})$
    if $\mathbf{v} \neq \mathbf{x} - \mathbf{w}$ **then**
      return accept;
    end if
  end for
  return reject;
end function

It has been noted in Section 3 that $\mathsf{solveCVP}$ calls $\mathsf{solveLWE}$ a total of n^{2c+2} times. So, $\mathsf{solveGapSVP}_{\zeta,\gamma}$ calls $\mathsf{solveLWE}$ a total of $N \cdot n^{2c+2}$ times.

We now consider the success probability of $\mathsf{solveGapSVP}_{\zeta,\gamma}$. As in Section 3, assume that $m = n^c$, $\alpha = 1/\sqrt{n}$ and $t = \mu_0/2$. The probability that a single call to $\mathsf{solveCVP}$ is successful is at least ε, where using (2), $\varepsilon = \left(1 - \exp\left(-n^c \exp(-2\pi\alpha^2)/8\right)\right)^{n^{2c+2}}$. The N calls to $\mathsf{solveCVP}$ in Algorithm $\mathsf{solveGapSVP}_{\zeta,\gamma}$ are independent. Let E be the event that all these calls are successful and so $\Pr[E] \geq \varepsilon^N$.

For $i = 1, \ldots, N$, let S_i be the event that the event $\mathbf{v} \neq \mathbf{x} - \mathbf{w}$ holds in the i-th iteration. The events $S_1, \ldots, S_N$ are independent (even when conditioned on E).

First consider the instance $(\mathbf{B}, r)$ to be NO instance of GapSVP$_{\zeta,\gamma}$. Let succNO be the event that algorithm $\mathsf{solveGapSVP}_{\zeta,\gamma}$ is successful on a NO instance. Then $\Pr[\mathsf{succNO}] = \Pr[\overline{S}_1 \wedge \cdots \wedge \overline{S}_N] \geq \Pr[\overline{S}_1 \wedge \cdots \wedge \overline{S}_N | E] \Pr[E] = \Pr[E] \cdot \left(\prod_{i=1}^N \Pr[\overline{S}_i | E]\right) \geq \varepsilon^N \cdot \left(\prod_{i=1}^N \Pr[\overline{S}_i | E]\right)$. It has been shown in [13] that $\Pr[\overline{S}_i | E] \approx 1$, $i = 1, \ldots, N$, and so we may assume that $\Pr[\mathsf{succNO}]$ is lower bounded by ε^N.

Next consider the instance $(\mathbf{B}, r)$ to be a YES instance of GapSVP$_{\zeta,\gamma}$. Let $\mathsf{succYES}$ be the event that algorithm $\mathsf{solveGapSVP}_{\zeta,\gamma}$ is successful on a YES instance. So, $\mathsf{succYES}$ is the event $S_1 \vee (\overline{S}_1 \wedge S_2) \vee \cdots \vee (\overline{S}_1 \wedge \cdots \wedge \overline{S}_{N-1} \wedge S_N)$.

For $i = 1, \ldots, N$, let δ be the common value of $\Pr[\overline{S}_i|E]$. It follows (using a probability calculation) that

$$\Pr[\mathsf{succYES}] \quad \geq \quad \Pr[\mathsf{succYES}|E]\Pr[E] = (1-\delta^N)\Pr[E] \geq (1-\delta^N)\varepsilon^N.$$

It has been shown in [13], that for a YES instance, $\delta = \Pr[\overline{S}_i|E] \leq 1 - 1/\mathsf{poly}(n)$. The $1 - 1/\mathsf{poly}(n)$ term arises from the asymptotic form of a result on statistical distance between the uniform distribution on $d' \cdot \mathcal{B}_n$ and the uniform distribution on $\mathbf{z} + d' \cdot \mathcal{B}_n$ with $\|\mathbf{z}\| \leq d$ and $d' = d \cdot \sqrt{n/(c \log n)}$. This result is proved in [10] the proof shows that the term $1 - 1/\mathsf{poly}(n)$ can be taken to be $1 - 3/n^2$. Using this we have $\delta \leq 1 - 3/n^2$. For simplicity, we take $\delta \leq 1 - 1/n^2$. So, $\Pr[\mathsf{succYES}] \geq (1 - (1 - 1/n^2)^N)\varepsilon^N$.

Between the NO and YES instances, the lower bound on the success probability is lesser for YES instances. As a result, the upper bound on the tightness gap for YES instances is higher and this upper bound is taken to be the upper bound on the overall tightness gap of the reduction. So, an upper bound on the tightness gap of the $\mathrm{GapSVP}_{\zeta,\gamma}$ to LWE reduction is

$$\left(N \cdot n^{2c+2}\right) / \left((1 - (1 - 1/n^2)^N)\varepsilon^N\right). \tag{5}$$

Following [10], for $N = n^2$, $(1 - (1 - 1/n^2)^N) \approx 1$ and so the tightness gap in (5) becomes

$$N \cdot n^{2c+2} \cdot \varepsilon^{-N} = n^{2c+4}\left(1 - \exp\left(-n^c \exp(-2\pi\alpha^2)/8\right)\right)^{-n^{2c+4}}. \tag{6}$$

We note that for $c = 1$, the expression in (6) is almost the same as the expression in (3). It has been shown in [16], that for $\alpha \leq 1/\sqrt{n}$, $\varepsilon \approx 1$ and so the tightness gap of $\mathrm{GapSVP}_{\zeta,\gamma}$ to $\mathrm{LWE}_{q,\Psi_\alpha}$ becomes

$$n^{2c+4}. \tag{7}$$

Remark: It is known [13] that for $\zeta(n) \geq 2^{n/2}$, the problem $\mathrm{GapSVP}_{\zeta,\gamma}$ is equivalent to the standard GapSVP_γ problem. The reduction from $\mathrm{GapSVP}_{\zeta,\gamma}$ to $\mathrm{LWE}_{q,\Psi_\alpha}$ given in [13] holds under the condition $q = q(n) \geq \zeta(n) \cdot \omega(\sqrt{\log n/n})$. So, for $q(n) \geq 2^{n/2} \cdot \omega(\sqrt{\log n/n})$, there is a classical reduction from GapSVP_γ to $\mathrm{LWE}_{q,\Psi_\alpha}$, where $\gamma = \gamma(n) \geq n/(\alpha\sqrt{\log n})$.

5 Reducing GapSVP_γ to Decision LWE

The remark at the end of Section 4 shows that there is a classical reduction of GapSVP_γ to $\mathrm{LWE}_{q,\Psi_\alpha}$ for $q(n) \geq 2^{n/2} \cdot \omega(\sqrt{\log n/n})$. So, if the modulus of the LWE problem is exponential in the dimension of the lattice, then the result from [13] provides a classical reduction of GapSVP_γ to LWE. A later work by Brakerski et al. [6] showed a reduction of GapSVP_γ to a decision version of LWE with polynomial sized modulus. The reduction is quite intricate and is built by composing reductions between several pairs of problems. The goal of the present section is to perform a concrete security analysis of the reduction provided in [6].

The LWE problem considered in Section 2.1 is a search problem. For the classical reduction of GapSVP_γ to LWE, decision versions of the LWE problem have been considered.

1. Let $\mathbf{s}$ be chosen uniformly at random from $\mathbb{Z}_q^n$. The $\mathrm{decLWE}_{n,m,q,\alpha}$ problem is to distinguish the uniform distribution over $\mathbb{T}_q^n \times \mathbb{T}$ from $A_{q,\mathbf{s},\alpha}$, where a list of m independent samples of the relevant distribution is provided as input.

2. Let $\mathbf{s}$ be chosen uniformly at random from $\{0,1\}^n$. The $\mathrm{binLWE}_{n,m,q,\alpha}$ problem is to distinguish the uniform distribution over $\mathbb{T}_q^n \times \mathbb{T}$ from $A_{q,\mathbf{s},\alpha}$, where a list of m independent samples of the relevant distribution is provided as input.

The difference between the decLWE and the binLWE problem lies in the method to select the secret $\mathbf{s}$. Given $n, q \geq 1$ and $\alpha \in (0,1)$, $\text{binLWE}_{n,m,q,\leq\alpha}$ is the problem which requires to solve $\text{binLWE}_{n,m,q,\beta}$ for any $\beta = \beta(\mathbf{s}) \leq \alpha$ [6].

Let $\mathcal{D}_0$ be the distribution $A_{q,\mathbf{s},\alpha}$ and $\mathcal{D}_1$ be the uniform distribution over $\mathbb{Z}_q^n \times \mathbb{T}$. For $i = 0, 1$, let $\mathcal{I} \xleftarrow{m} \mathcal{D}_i$ denote the selection of a list $\mathcal{I}$ of m independent samples from $\mathcal{D}_i$. Let $\mathcal{A}$ be a distinguisher for $\text{decLWE}_{n,m,q,\phi}$. Let $\mathcal{A}(\mathcal{I}) \Rightarrow 1$ denote the event that $\mathcal{A}$ produces 1 as output. The advantage of $\mathcal{A}$ is the following.

$$\mathsf{Adv}(\mathcal{A}) \quad = \quad |\Pr[\mathcal{A}(\mathcal{I}) \Rightarrow 1 : I \xleftarrow{m} \mathcal{D}_0] - \Pr[\mathcal{A}(\mathcal{I}) \Rightarrow 1 : I \xleftarrow{m} \mathcal{D}_1]|. \tag{8}$$

Similarly, one defines the advantage of a distinguisher for $\text{binLWE}_{n,m,q,\phi}$.

The classical reduction in [6] reduces GapSVP to binLWE. This reduction is done in several steps. The first step is Peikert's reduction of GapSVP to LWE with exponential size modulus. The goal of the following steps is to reduce the LWE problem with exponential size modulus to binLWE problem with polynomial size modulus. A trade-off is an increase in the dimension. The various steps of the overall reduction are as follows.

Reducing GapSVP_γ to $\text{LWE}_{k,m_1,q_1,\alpha_1}$: This follows from Peikert's result [13]. Here $\alpha_1 \in (0,1)$, $q_1 \geq 2^{k/2} \cdot \omega(\sqrt{\log k/k})$, $\gamma \geq k/(\alpha_1\sqrt{\log k})$ and $m_1 = k^c$ for some constant $c \geq 1$. For simplicity, in the following, we will assume $q_1 = 2^{k/2}$.

Suppose W_0 is an algorithm to solve $\text{LWE}_{k,m_1,q_1,\alpha_1}$. Then following the analysis in Section 4, there is an algorithm W to solve GapSVP_γ where the number of times W calls W_0 is k^{2c+4} (which is obtained from (7) by replacing n with k).

Reducing $\text{LWE}_{k,m_1,q_1,\alpha_1}$ to $\text{decLWE}_{k,m_1,q_1,\alpha_2}$: This follows as a special case of Theorem 3.1 in [12]. Here $1/q_1 < \alpha_1 < 1/\omega(\sqrt{\log n})$ and $\alpha_2 = \alpha_1 \cdot \omega(\log k)$.

To determine the tightness gap of the reduction, we follow the proof of Theorem 3.1 in the case where $q_1 = 2^{k/2}$. Let W_1 be an algorithm to solve $\text{decLWE}_{k,m_1,q_1,\alpha_2}$. The proof of Theorem 3.1 in [12] uses W_1 to first construct an algorithm W_1' following the construction used in Lemma 4.1 of [15]. Specifically, Lemma 4.1 of [15] shows how to boost the advantage of a distinguisher for the distributions $A_{q_1,\mathbf{s},\chi}$ and $U(\mathbb{Z}_{q_1}^n \times \mathbb{Z}_{q_1})$. The same method can be used to boost the advantage of a distinguisher for the distributions $A_{q_1,\mathbf{s},\alpha_2}$ and the uniform distribution on $\mathbb{Z}_{q_1}^n \times \mathbb{T}$. This is the situation considered in Theorem 3.1 of [12].

Let ζ_1 be the advantage of W_1 and c_1 and c_2 be such that W_1 is successful on a fraction k^{-c_1} of all possible secrets and

$$\zeta_1 \quad \geq \quad k^{-c_2}. \tag{9}$$

Following the method of Lemma 4.1 in [15] it is possible to construct W_1' which accepts with probability exponentially close to one on inputs from $A_{q_1,\mathbf{s},\alpha_2}$ and rejects with probability exponentially close to one on inputs from the uniform distribution over $\mathbb{Z}_{q_1}^n \times \mathbb{T}$. From the proof of Lemma 4.1 in [15] we have that the algorithm W_1' calls the algorithm W_1 a total of $k^{c_1+2c_2+2}$ times.

The proof of Theorem 3.1 in [12] uses W_1' to construct an algorithm W_0 which solves $\text{LWE}_{k,m_1,q_1,\alpha_1}$. The secret $\mathbf{s} = (s_1, \ldots, s_k)$. The components $s_1, \ldots, s_k$ are determined one by one. Consider the determination of s_1. This is determined iteratively as $s_1 \bmod 2$, followed by $s_1 \bmod 2^2$, followed by $s_1 \bmod 2^3$, up to at most $s_1 \bmod 2^{k/2}$. Given the value of $s_1 \bmod 2^i$, there are only two possible values for $s_1 \bmod 2^{i+1}$. A single call to W_1' can be used to determine the correct value. So, to find s_1, at most $k/2$ calls to W_1' are required, and to find the entire vector $\mathbf{s}$, at most $k^2/2$ calls to W_1' are required. Each call to W_1' requires $k^{c_1+2c_2+2}$ calls to W_1. So, the number of times W_0 calls W_1 is

$$k^{c_1+2c_2+4}. \tag{10}$$

Reducing decLWE$_{k,m_1,q_1,\alpha_2}$ to binLWE$_{n,m_1,q_1,\leq\sqrt{10n}\alpha_2}$: This reduction follows from Theorem 4.1 of [6]. Here $n \geq (k+1)\log_2 q_1 + 2\log_2(1/\delta)$, $\alpha_2 \geq \sqrt{\ln(2n(1+1/\varepsilon_1))/\pi}/q_1$, where $\delta > 0$ and $\varepsilon_1 \in (0, 1/2)$. Suppose there is an algorithm W_2 for binLWE$_{n,m_1,q_1,\leq\sqrt{10n}\alpha_2}$ which has advantage ζ_2. Theorem 4.1 of [6] shows an algorithm W_1 for decLWE$_{k,m_1,q_1,\alpha_2}$ with advantage ζ_1 where

$$\zeta_1 \quad \geq \quad \frac{\zeta_2 - \delta}{3m_1} - \frac{41\varepsilon_1}{2} - 2^{-k-1}. \tag{11}$$

We assume that W_1 calls W_2 once.

Remark: We note a peculiarity in (11). The number of samples m_1 appears in the denominator of the right hand side. So, as m_1 increases, the right hand side decreases. In other words, as the number of samples increases, the lower bound on the advantage ζ_1 decreases. Intuitively, one may expect that as the number of samples increases, more information is obtained, and so the advantage should be non-decreasing. This, however, does not hold for (11) indicating a non-intuitive property of the reduction.

Reducing binLWE$_{n,m_1,q_1,\leq\sqrt{10n}\alpha_2}$ to binLWE$_{n,m_1,q_2,\leq\alpha_3}$: This reduction follows from Corollary 3.2[2] of [6]. Here $q_1 \geq q_2 \geq \sqrt{2\ln(2n(1+1/\varepsilon_2))} \cdot (\sqrt{n}/\alpha_2)$ and $\alpha_3^2 \geq 10n\alpha_2^2 + (4n/(\pi q_2^2))\ln(2n(1+1/\varepsilon_2))$ where $\varepsilon_2 \in (0, 1/2)$. Suppose there is an algorithm W_3 for binLWE$_{n,m_1,q_2,\leq\alpha_3}$ having advantage ζ_3. Corollary 3.2 of [6] shows an algorithm W_2 for binLWE$_{n,m_1,q_1,\leq\sqrt{10n}\alpha_2}$ with advantage ζ_2 where

$$\zeta_2 \quad \geq \quad \zeta_3 - 14\varepsilon_2 m_1. \tag{12}$$

We will assume that W_2 calls W_3 once.

Reducing binLWE$_{n,m_1,q_2,\leq\alpha_3}$ to binLWE$_{n,m_2,q_2,\alpha_3}$: This reduction follows from Lemma 2.15 of [6]. Suppose there is an algorithm W_4 for binLWE$_{n,m_2,q_2,\alpha_3}$ having advantage ζ_4. Lemma 2.15 of [6] states that the algorithm W_3 for binLWE$_{n,m_1,q_2,\leq\alpha_3}$ has advantage ζ_3 where $\zeta_3 \geq 1/3$. Further, both m_1 and the number of times W_3 calls W_4 is $\mathfrak{p}$ where

$$\mathfrak{p} \quad = \quad \text{poly}(m_2, 1/\zeta_4, n, \log q_2). \tag{13}$$

We assume that there are constants $d_1, d_2 > 0$, such that $m_2 = n^{d_1}$ and $\zeta_4 \geq n^{-d_2}$.

Putting together the various reductions, yields a reduction from GapSVP$_\gamma$ on a lattice of dimension k to binLWE$_{n,m_2,q_2,\alpha_3}$. The number of times the algorithm W_4 (for solving binLWE$_{n,m_2,q_2,\alpha_3}$) is called by the algorithm W (for solving GapSVP$_\gamma$) is obtained from the above analysis to be the following.

$$k^{2c+4} \cdot k^{c_1+2c_2+4} \cdot \mathfrak{p}. \tag{14}$$

The relations among the various parameters are as follows.

1. $\gamma \geq k/(\alpha_1\sqrt{\log k})$;
2. $q_1 = 2^{k/2}$;
3. $m_1 = k^c$ for some constant $c \geq 1$;
4. $1/q_1 < \alpha_1 < 1/\omega(\sqrt{\log n})$ and $\alpha_2 = \alpha_1 \cdot \omega(\log k)$;

[2] A distribution $\mathcal{D}$ over $\mathbb{Z}^n$ is (B, δ)-bounded, for $B, \delta \in \mathbb{R}$, if the probability that $\mathbf{x} \leftarrow \mathcal{D}$ has norm greater than B is at most δ. Corollary 3.2 of [6] is stated in terms of (B, δ) distribution $\mathcal{D}$. In the present context, $\mathcal{D}$ is the uniform distribution over $\{0, 1\}$ which is $(\sqrt{n}, 0)$-bounded.

5. The constants c_1 and c_2 are such that W_1 is successful on a fraction k^{-c_1} of all possible secrets and $\zeta_1 \geq k^{-c_2}$;
6. $n \geq (k+1)\log_2 q_1 + 2\log_2(1/\delta)$;
7. $\alpha_2 \geq \sqrt{\ln(2n(1+1/\varepsilon_1))/\pi}/q_1$, and $\zeta_1 \geq \frac{\zeta_2-\delta}{3m_1} - \frac{41\varepsilon_1}{2} - 2^{-k-1}$, where $\delta > 0$ and $\varepsilon_1 \in (0, 1/2)$;
8. $q_1 \geq q_2 \geq \sqrt{2\ln(2n(1+1/\varepsilon_2))} \cdot (\sqrt{n}/\alpha_2)$, $\alpha_3^2 \geq 10n\alpha_2^2 + (4n/(\pi q_2^2))\ln(2n(1+1/\varepsilon_2))$, and $\zeta_2 \geq \zeta_3 - 14\varepsilon_2 m_1$, where $\varepsilon_2 \in (0, 1/2)$;
9. $\zeta_3 \geq 1/3$;
10. $m_1 = \mathfrak{p} = \mathrm{poly}(m_2, 1/\zeta_4, n, \log q_2)$;
11. $m_2 = n^{d_2}$ and $\zeta_4 \geq n^{-d_1}$ for constant $d_1, d_2 > 0$.

Note that

$$\begin{aligned}
\zeta_1 &\geq \frac{\zeta_2-\delta}{3m_1} - \frac{41\varepsilon_1}{2} - 2^{-k-1} \geq \frac{\zeta_3}{3m_1} - \frac{14\varepsilon_2}{3} - \frac{\delta}{3m_1} - \frac{41\varepsilon_1}{2} \geq \frac{1}{9m_1} - \frac{14\varepsilon_2}{3} - \frac{\delta}{3m_1} - \frac{41\varepsilon_1}{2}, \\
\alpha_3^2 &\geq 10n\alpha_2^2 + \frac{4n}{\pi q_2^2}\ln(2n(1+1/\varepsilon_2)) \geq 10n\alpha_1^2\omega(\log^2 k) + \frac{4n}{\pi q_2^2}\ln(2n(1+1/\varepsilon_2)).
\end{aligned}$$

Performing a meaningful concrete security analysis with the exact form of the above relations is almost impossible. To simplify the analysis, we ignore logarithmic factors. Also, we will assume that the parameters ε_1, ε_2 and δ can be chosen in a manner (say, $1/\mathrm{poly}(n)$) such that they do not have much effect on the concrete security analysis. Using these and other reasonable simplifications, we have the following relations.

$$\begin{aligned}
&q_1 = 2^{k/2}; \quad n = k^2; \\
&\alpha_1 = \alpha_2 = \alpha_3/\sqrt{n} = \alpha_3/k; \\
&\gamma = k/\alpha_1 = k^2/\alpha_3; \\
&c_1 = c_2; \\
&k^{-c_2} = \zeta_1 = 1/m_1 = k^{-c}, \\
&q_2 = \sqrt{n}/\alpha_2 = n/\alpha_3; \\
&k^c = m_1 = \mathfrak{p} = 1/\zeta_4 = n^{d_1}.
\end{aligned} \tag{15}$$

From (15), we have $c_1 = c_2 = c = 2d_1$. Using these in (14), the overall tightness gap is obtained to be

$$n^{6d_1+4}. \tag{16}$$

The tightness gap given by (16) is to be compared to the tightness gap of Regev's reduction given in (4). For practical scenarios (see below for a numerical example), the tightness gap given by (16) is larger than that given by (4).

Summary: We have the following concrete form of the reduction of GapSVP to binLWE.

> If there is an algorithm which solves $\mathrm{binLWE}_{n,m_2,q_2,\alpha_3}$, with $q_2 = n/\alpha_3$ having advantage $\zeta_4 = n^{-d_1}$, then there is an algorithm to solve $\mathrm{GapSVP}_{k^2/\alpha_3}$ on a lattice of dimension $k = \sqrt{n}$. The tightness gap of the reduction is given by n^{6d_1+4}.

The problem $\mathrm{binLWE}_{n,m_2,q_2,\alpha_3}$ would be used as a basis for proving security of cryptosystems. We consider $\alpha_3 = 1/\sqrt{n} = 1/k$.

As a numerical example, consider $n = 2^{10}$. Aiming at 128-bit security, ζ_4 would be 2^{-128} and so for $n = 2^{10}$, $d_1 = 12.8$. The tightness gap in (16) is then about 2^{808}. In other words, the quantitative effect of the reduction is

the following. If T is the time required to solve $\text{binLWE}_{n,m_2,q_2,\alpha_3}$ with advantage 2^{-128} on a lattice of dimension 2^{10}, then there is an algorithm to solve GapSVP_γ for a lattice of dimension $k = \sqrt{n} = 2^5$ and $\gamma = k^3 = 2^{15}$ which takes time $2^{808}T$. So, the tightness gap is 2^{808}. In comparison, for $n = 2^{10}$ and 128-bit security, the tightness gap in [7, 16] has been obtained to be 2^{524}.

Note that the dimension of the lattice for which GapSVP is to be solved is $\sqrt{n}$ where n is the dimension of the lattice for which binLWE is to be solved. Brakerski et al. [6] mention this point. Due to the drawback of the quadratic loss in the dimension, they mention as an open problem the task of obtaining a reduction where such a quadratic loss does not occur. In their words, this would constitute a "full dequantization" of Regev's reduction.

The issue of tightness gap has not been considered in [6]. For the GapSVP to binLWE reduction to be meaningfully used to derive parameters for practical cryptosystems, the tightness gap needs to be taken into consideration. So, for a full dequantization of Regev's reduction which can also be used in practice, one needs a *tight* reduction which does not suffer the quadratic loss in the dimension.

6 Conclusion

We have performed a concrete security analysis of the tightness gap in the classical reduction of the shortest vector problem to the LWE problem given by Brakerski et al. [6]. Previous works [7, 16] had already pointed out that the tightness gap in the quantum reduction of Regev [14] is huge. Our analysis shows that the tightness gap of the classical reduction by Brakerski et al. is more than that of Regev's original quantum reduction. This leaves open the question of obtaining a tight reduction of a worst case lattice problem to LWE, or, showing that there is no such reduction.

References

[1] Erdem Alkim, Roberto Avanzi, Joppe Bos, Leo Ducas, Antonio de la Piedra, Thomas Poppelmann, Peter Schwabe, Douglas Stebila, Martin R. Albrecht, Emmanuela Orsini, Valery Osheter, Kenneth G. Paterson, Guy Peer, and Nigel P. Smart. NewHope: algorithm specifications and supporting documentation. https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions, 2019.

[2] Erdem Alkim, Joppe Bos, Leo Ducas, Patrick Longa, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Christopher Peikert, Ananth Raghunathan, and Douglas Stebila. FrodoKEM: Learning With Errors key encapsulation. https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions, 2019.

[3] Roberto Avanzi, Joppe Bos, Leo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Kyber: algorithm specifications and supporting documentation. https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions, 2009.

[4] Hayo Baan, Sauvik Bhattacharya, Scott Fluhrer, Oscar Garcia-Morchon, Thijs Laarhoven, Rachel Player, Ronald Rietman, Markku-Juhani O. Saarinen, Ludo Tolhuizen, Jose-Luis Torre-Arce, and Zhenfei Zhang. Round5: KEM and PKE based on (Ring) Learning With Rounding. https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions, 2019.

[5] Daniel J. Bernstein. Comparing proofs of security for lattice-based encryption. Cryptology ePrint Archive, Report 2019/691, 2019. https://eprint.iacr.org/2019/691.

[6] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on*

Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013, pages 575–584. ACM, 2013.

[7] Sanjit Chatterjee, Neal Koblitz, Alfred Menezes, and Palash Sarkar. Another look at tightness II: practical issues in cryptography. In Raphael C.-W. Phan and Moti Yung, editors, *Paradigms in Cryptology - Mycrypt 2016. Malicious and Exploratory Cryptology - Second International Conference, Mycrypt 2016, Kuala Lumpur, Malaysia, December 1-2, 2016, Revised Selected Papers*, volume 10311 of *Lecture Notes in Computer Science*, pages 21–55. Springer, 2016.

[8] Jan-Pieter DAnvers, Angshuman Karmakar, Sujoy Sinha Roy, and Frederik Vercauteren. SABER: Mod-LWR based KEM (round 2 submission). https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions, 2019.

[9] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Cynthia Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 197–206. ACM, 2008.

[10] Oded Goldreich and Shafi Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000.

[11] Xianhui Lu, Yamin Liu, Dingding Jia, Haiyang Xue, Jingnan He, Zhenfei Zhang, Zhe Liu, Hao Yang, Bao Li, and Kunpeng Wang. LAC: Lattice-based Cryptosystems. https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions, 2019.

[12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. Cryptology ePrint Archive, Report 2011/501, https://eprint.iacr.org/2011/501, 2011. An abridged version of this paper appeared in the proceedings of Eurocrypt 2012.

[13] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 333–342. ACM, 2009.

[14] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93. ACM, 2005.

[15] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, 2009.

[16] Palash Sarkar and Subhadip Singha. Verifying solutions to LWE with implications for concrete security. *Advances in Mathematics of Communications*, 2020. https://www.aimsciences.org/article/doi/10.3934/amc.2020057.