

Concrete Analysis of Approximate Ideal-SIVP to Decision Ring-LWE Reduction

Neal Koblitz

Department of Mathematics, Box 354350,
University of Washington, Seattle, WA 98195 U.S.A.
email: koblitz@uw.edu

Subhabrata Samajder

Department of Computer Science and Engineering,
Indraprastha Institute of Information Technology, Delhi, 110020 India
email: subhabrata@iiit.ac.in

Palash Sarkar and Subhadip Singha

Applied Statistics Unit,
Indian Statistical Institute, Kolkata, India 700108
email: {palash, subha.r}@isical.ac.in

March 5, 2022

Abstract

A seminal 2013 paper by Lyubashevsky, Peikert, and Regev proposed basing post-quantum cryptography on ideal lattices and supported this proposal by giving a polynomial-time security reduction from the approximate Shortest Independent Vectors Problem (SIVP) to the Decision Learning With Errors (DLWE) problem in ideal lattices. We give a concrete analysis of this multi-step reduction. We find that the tightness gap in the reduction is so great as to vitiate any meaningful security guarantee, and we find reasons to doubt the feasibility in the foreseeable future of the quantum part of the reduction. In addition, when we make the reduction concrete it appears that the approximation factor in the SIVP problem is far larger than expected, a circumstance that causes the corresponding approximate-SIVP problem most likely not to be hard for proposed cryptosystem parameters. We also discuss implications for systems such as Kyber and SABER that are based on module-DLWE.

Keywords: ideal lattices, shortest vector problem, ring learning with errors, concrete analysis.

Mathematics Subject Classification: 94A60

Contents

1 Introduction	2
2 Preliminaries	5
3 Reducing K-SIVP$_{\gamma}$ to search ring-LWE$_{q, \leq \alpha}$	8
4 Reducing search ring-LWE to ring-DLWE	13

1	5	The tightness gap in the K-SIVP$_{\gamma}$ to ring-DLWE$_{q,\tau}$ reduction	22
2	6	Problems with the quantum part of the reduction	24
3	7	Subsequent work	28
4	8	Conclusion	30
5	A	Details of the parameters of \mathcal{A}_2 and \mathcal{A}_3	34
6	B	Details of the analysis in Section 4.4	35

7 1 Introduction

8 In 2013 Lyubashevsky, Peikert, and Regev [19] published a security reduction in support of proposed post-
9 quantum cryptography based on the difficulty of the Decision Learning With Errors (DLWE) problem in a
10 lattice coming from the embeddings of an ideal of a number field¹. Their elaborate, multi-step reduction showed
11 that the worst-case γ -approximate Shortest Independent Vectors Problem (SIVP $_{\gamma}$) for ideal lattices could be
12 solved with polynomially many calls to an oracle that solves DLWE for ideal lattices. Our purpose in this paper
13 is to analyze this reduction in concrete terms.

14 The U.S. government’s NIST is currently running a multi-year competition² to select candidates for standardi-
15 sation of post-quantum public key cryptography. Some of the proposals under consideration are based on lattices.
16 Two of the lattice-based finalists, namely Kyber [7] and SABER [12], are based on module lattices [18, 25] which
17 are generalisations of ideal lattices. In [25] Peikert and Pepin show that ring-LWE reduces to module-LWE for
18 a given size of the module and ring and then argue for the security of module-LWE-based systems by citing
19 presumed hardness of ring-LWE. In addition, in [18] Langlois and Stehlé give a reduction from approximate
20 module-SIVP to module-DLWE. In §7.1 we discuss the generalisation from ring-DLWE to module-DLWE.

21 1.1 The structure of the approximate ideal-SIVP to decision ring-LWE reduction

22 The structure of the reduction in [19] – a nested sequence of intermediate reductions – gives rise to two difficulties
23 from a practice-oriented perspective. In the first place, the tightness gaps multiply from one reduction to the
24 next. If algorithm A calls on algorithm B m times, and B calls on C n times, then there are mn calls on C .
25 We found that the cumulative tightness gap in the reduction is so great as to render the security guarantee
26 meaningless for practical parameter values.

27 In the second place, seven of the nested reductions take place within a quantum computer. We have reasons
28 to doubt the feasibility of the quantum part of the reduction even assuming the advent of quantum computers
29 that are scaled to much larger size than what is needed to break RSA and ECC.
30

31 1.2 Restricting to a special class of lattices

32 From the beginning, the strongest argument advanced for lattice-based cryptography has been worst-to-average
33 case reductions. The hardest instances of problems such as the approximate Shortest Vector Problem (SVP)
34 and the approximate Shortest Independent Vectors Problem (SIVP) can sometimes be shown to reduce to a
35 random instance of the lattice problem that the proposed cryptosystem is based on. This argument loses validity

¹Earlier D. Stehlé *et al.* [34] had published a security reduction for ideal lattices, but their reduction only goes as far as search ring-LWE, not ring-DLWE.

²<https://csrc.nist.gov/Projects/post-quantum-cryptography/round-3-submissions>, accessed on February 8, 2022.

1 if lattices are chosen from a special class, such as the class of ideal lattices, rather than from the set of all general
 2 lattices, unless one has evidence that the shortest vector problems and their approximate variants do not lose
 3 any of their worst-case intractability when restricted to the special class of ideal lattices.

4 As far as we are aware, no such evidence exists. After arguing for many years that worst-to-average case
 5 reductions are important in order to have confidence in security, it seems that some promoters of lattice-based
 6 systems have undermined that argument by changing course and now preferring to work in a special subclass of
 7 lattices.

8 From a number theory perspective, one reason to wonder about the effect of specializing to ideal lattices is that
 9 they have much more structure than general lattices, notably the presence of isomorphisms between different
 10 embeddings of the number field. In the cyclotomic case (and, more generally, in the case of non-cyclotomic
 11 Galois fields as well) the isomorphisms are all automorphisms (which permute the roots of unity), and those
 12 automorphisms were in fact used to good effect in the security reductions in [19].

13 But the much greater structure and symmetry, especially in the cyclotomic case, also make it likely that the
 14 supposedly hard shortest vector problems that are the basis for the security of lattice-based protocols are in fact
 15 much easier in this restricted setting. For example, as pointed out in [19], in cyclotomic ideal lattices SIVP_γ is
 16 trivially equivalent to SVP_γ , because a short vector can be multiplied by roots of unity to get an entire basis of
 17 vectors of the same length. In contrast, in the general case only a much weaker result is known, namely, that
 18 $\text{SIVP}_{\sqrt{n}\gamma}$ reduces in polynomial time to SVP_γ , where n is the dimension of the lattice [20]. This suggests that
 19 for general lattices SIVP and SIVP_γ are strictly harder than SVP and SVP_γ – a separation that disappears when
 20 restricted to cyclotomic ideal lattices. Thus, even if SVP and approximate SVP for cyclotomic ideal lattices were
 21 to be as hard as for general lattices, SIVP and approximate SIVP for cyclotomic ideal lattices would likely be
 22 easier.

23 In the simplest case of cyclotomic fields generated by m -th roots of unity with $m = 3, 4$, i.e., with $n = \varphi(m) =$
 24 2 , SVP/SIVP is trivial for any ideal lattice³, whereas the general SVP/SIVP in two dimensions is not completely
 25 trivial. It is not yet clear whether this gap in difficulty increases in higher dimensions, but there is no reason to
 26 assume that it does not. Moreover, for the full lattice R , where R is the ring of integers of a cyclotomic field,
 27 Lemma 2.9 of [19] tells us that the ℓ_2 -norm of a shortest vector is $\geq \sqrt{n}$. Since $\|1\| = \sqrt{n}$, that means that 1 is a
 28 shortest vector. If the ideal \mathcal{I} is principal (which is true for all ideals for $n = 2, 4, 8, 16$), then for each imbedding
 29 σ the image of \mathcal{I} is just a scaled (and rotated) version of the image of R . That does not immediately lead to a
 30 simple result for the shortest vector in \mathcal{I} , but it certainly suggests a close relationship between the geometry of
 31 \mathcal{I} and the geometry of the unit ideal lattice R , for which SVP and SIVP are trivial.

32 1.3 The role of the approximation factor

33 Another reason why the reduction of SIVP_γ to SVP_γ for cyclotomic ideal lattices is troubling is that for $\gamma > n$,
 34 Goldreich and Goldwasser [14] showed that SVP_γ is unlikely to be NP-hard. In [19] the problem assumed to be
 35 hard is SIVP_γ where $\gamma = \tilde{O}(\sqrt{n}/\alpha)$ in which $\alpha < (\ln n/n)^{1/2}$ is much less than 1. The Goldreich-Goldwasser
 36 result shows that this approximate SIVP problem for cyclotomic ideal lattices is almost certainly not NP-hard.
 37 Further, the GapSVP problem on ideal lattices is easy [30], while it is conjectured to be hard for general lattices.

38 The last sentence in the statement of the “main theorem” (Theorem 3.6) of [19] says that the target problem
 39 for the reduction from $\tilde{O}(\sqrt{n}/\alpha)$ -approximate SIVP can be taken to be decision ring-LWE $_{q, D_{r_0}}$ with

$$r_0 = \alpha(n\ell/\log(n\ell))^{1/4}, \tag{1}$$

40 where ℓ is the number of queries made by the DLWE-distinguisher. Based on the estimate for α in the previous
 41 paragraph, this suggests a $\tilde{O}(n)$ approximation factor for the short independent vectors problems that are the
 42 basis for the security of DLWE-based protocols. This appears to be incorrect.

³The case $m = 3$ corresponds to tiling of the plane using equilateral triangles, whereas the case $m = 4$ corresponds to tiling the plane with squares.

1 First of all, cryptosystems based on the decision-LWE problem generally have a fixed and publicly known
 2 Gaussian distribution width r_0 , which is assumed to be roughly of order $n^{-1/2}$. In that case, ignoring the
 3 contribution of ℓ , from (1) we find a larger approximation factor $\tilde{O}(n^{5/4})$.

4 In the second place, when filling in the details of the reduction, we found that in place of (1) we needed

$$r_0 = \alpha(N_2 n \ell / \log(N_2 n \ell))^{1/4}, \quad (2)$$

5 where N_2 is a parameter for the reduction algorithm that is of order at least $n^2 \delta_2^{-2}$. This causes the approximation
 6 factor in the supposedly intractable SIVP to be greater than $n^{7/4} \delta_2^{-1/2}$. There is little reason to have confidence
 7 in the intractability of SIVP $_{n^{7/4} \delta_2^{-1/2}}$ for cyclotomic ideal lattices.

8 1.4 Efficiency versus security

9 Increased efficiency of implementation is the main reason for specializing lattice-based cryptography to the
 10 lattices coming from ideals of number rings and, in particular, to cyclotomic ideal lattices. As explained in [19],
 11 general lattices “are rather inefficient due to an inherent quadratic overhead in the use of LWE,” whereas ideal
 12 lattices provide a major speed-up in running time and a reduction in the size of public keys by a factor of n , the
 13 dimension of the lattice.

14 The abstract of [19] also uses the same term *efficient* in a very different sense of the word in describing the
 15 security reductions in their paper. When they speak of an “efficient security reduction” they mean a polynomial-
 16 time quantum reduction. The question of practical feasibility is not addressed in the paper.

17 This is an unfortunate omission. Ever since Bellare, Rogaway, and others argued for “practice-oriented
 18 provable security” in the 1990s [3], it has been widely recognized that a close examination of security reductions to
 19 determine the real-world guarantees that they give is essential. Moreover, it is a cardinal principle of cryptography
 20 that efficiency of usage should not be prioritized over meaningful evidence of security.

21 In nearly a decade since the appearance of [19] several attacks on ideal lattice problems have confirmed the
 22 intuition that such lattices are more vulnerable to attacks, both classical and quantum, than general lattices [11,
 23 5]. These works suggest that from a security standpoint cryptography based on ring-LWE and similar ideal
 24 lattice problems might not stand the test of time.

25 1.5 Related reductions

26 In [29] Regev introduced the LWE problem and gave a reduction from approximate SIVP over general lattices
 27 to Decision LWE. This work is generally considered a breakthrough in lattice-based cryptography and spurred
 28 a great deal of subsequent research. The concrete aspect of the reduction in [29] was analysed in [10] and the
 29 analysis was refined in [33]. Commenting on the concrete analysis of [29] in [10], Bernstein [4] remarked that “the
 30 loss of tightness is gigantic.” A different concrete analysis of the reduction in [29] was carried out in [13] and this
 31 work also considered increasing the value of n to compensate for the tightness gap. We revisit the tightness gap
 32 of the reduction in [29] and obtain a more accurate estimate. We point out several aspects that were overlooked
 33 or were overestimated in [10, 33, 13]. The resulting gap turns out to be greater than the previous estimates. Our
 34 present work shows that the tightness gap in the reduction in [19] is even greater. Further, we argue that trying
 35 to increase the value of n to compensate for the tightness gap in [29, 19] is not a meaningful exercise because of
 36 the very large values of n that would be needed.

37 The reduction in [29] is quantum. Later work pursued the goal of obtaining a classical reduction. The first
 38 result in this direction was obtained by Peikert [23], who gave a classical reduction from GapSVP over general
 39 lattices to the LWE problem. The drawback of this result was that it required an exponential size modulus.
 40 A subsequent work by Brakerski et al. [8] gave a reduction where the modulus is of polynomial size. This
 41 reduction was also unsatisfactory since it reduced GapSVP on a lattice of dimension \sqrt{n} to the n -dimensional

1 LWE problem. Since GapSVP is not hard over ideal lattices, the approach adopted in [23, 8] is not meaningful
 2 for such lattices. A concrete analysis of the reductions in [23, 8] was carried out in [32].

3 1.6 Outline of the paper

4 The background and preliminaries required for the paper are given in §2. The reduction in [19] can be divided into
 5 two parts. The first part is a reduction from approximate ideal-SIVP to the search ring-LWE problem, while the
 6 second part is a reduction from the search ring-LWE problem to the decision ring-LWE problem. The concrete
 7 analysis of the first part is described in §3 and that for the second part is described in §4. The two parts are
 8 combined and the end-to-end reduction from approximate ideal-SIVP to decision ring-LWE is summarised in §5.
 9 A detailed discussion of several problematic issues with the quantum aspect of the reduction is given in §6. The
 10 reduction in [19] holds for cyclotomic number fields. A follow-up work [26] gave a reduction from approximate
 11 ideal-SIVP to decision ring-LWE for any number field. In §7 we perform a concrete analysis of the reduction
 12 in [26] and show that the tightness gap is much larger than the tightness gaps of either [29] or [19]. Section 8
 13 has concluding remarks. Some important mathematical details of the reduction are presented in Appendices A
 14 and B.

15 2 Preliminaries

16 A brief summary of the relevant concepts is provided below. For further details the reader may refer to [19]. We
 17 note though that at certain places we have simplified the description that is in [19].

18 By \mathbb{Z} , \mathbb{Q} , \mathbb{R} and \mathbb{C} we will denote the sets of integers, rationals, reals and complex numbers respectively.
 19 Let n be a positive integer. For a vector $\mathbf{a} = (a_1, \dots, a_n)$ in \mathbb{R}^n or \mathbb{C}^n , the ℓ_2 -norm of \mathbf{a} is defined to be
 20 $\|\mathbf{a}\| = (|a_1|^2 + \dots + |a_n|^2)^{1/2}$ and the ℓ_∞ -norm of \mathbf{a} is defined to be $\|\mathbf{a}\|_\infty = \max_{i \in [n]} |a_i|$. We will mostly consider
 21 the ℓ_2 -norm. At a few places, the ℓ_∞ -norm is used and we will explicitly identify these cases.

22 Let s_1 and s_2 be non-negative integers such that $s_1 + 2s_2 = n$. The space $H \subset \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ is defined as

$$H = \{(x_1, \dots, x_n) \in \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2} : x_{s_1+s_2+j} = \overline{x_{s_1+j}}, j = 1, \dots, s_2\}. \quad (3)$$

23 Using the inner product on H induced on it by \mathbb{C}^n , it can be shown that H is isomorphic to \mathbb{R}^n as an inner product
 24 space. For $j \in [n]$, let $\mathbf{e}_j \in \mathbb{C}^n$ be the vector which has 1 in its j -th component and 0 elsewhere. An orthonormal
 25 basis for H is given by $\{\mathbf{h}_i\}_{i \in [n]}$, where for $j \in [s_1]$, $\mathbf{h}_j = \mathbf{e}_j$ and for $s_1 < j \leq s_1 + s_2$, $\mathbf{h}_j = (\mathbf{e}_j + \mathbf{e}_{j+s_2})/\sqrt{2}$,
 26 $\mathbf{h}_{j+s_2} = \sqrt{-1}(\mathbf{e}_j - \mathbf{e}_{j+s_2})/\sqrt{2}$. When $\mathbf{x} \in H$ is written in terms of the orthonormal basis as $\mathbf{x} = \sum_{i=1}^n a_i \mathbf{h}_i$ with
 27 $(a_1, \dots, a_n) \in \mathbb{R}^n$, the norm of \mathbf{x} is simply $\|(a_1, \dots, a_n)\|$.

28 A lattice is a discrete additive subgroup of H . Let $B = \{\mathbf{b}_1, \dots, \mathbf{b}_n\} \subset H$ be a set of linearly independent
 29 vectors. The full rank lattice generated by B is defined to be $\mathcal{L}(B) = \{\sum_{i=1}^n z_i \mathbf{b}_i : (z_1, \dots, z_n) \in \mathbb{Z}^n\}$. Given
 30 a lattice basis B , the fundamental parallelepiped $\mathcal{P}(B)$ is defined to be the set $\{B\mathbf{x} : \mathbf{x} \in \mathbb{R}^n, 0 \leq x_i < 1\}$.
 31 For an n -dimensional lattice Λ , let $\lambda_i(\Lambda)$ with $i \in \{1, \dots, n\}$, be the least real number r such that Λ has
 32 i linearly independent vectors with the longest having length r with respect to the ℓ_2 -norm. In particular,
 33 $\lambda_1(\Lambda) = \min_{\mathbf{x} \in \Lambda \setminus \{0\}} \|\mathbf{x}\|$ is called the minimum distance of the lattice. The dual of a lattice Λ is defined to be
 34 $\Lambda^* = \{\mathbf{x} \in H : \langle \mathbf{y}, \mathbf{x} \rangle \in \mathbb{Z}, \text{ for all } \mathbf{y} \in \Lambda\}$, where $\langle \mathbf{y}, \mathbf{x} \rangle = \sum y_i \bar{x}_i$ is the inner product. Theorem 2.1 of [2] shows
 35 that for any n -dimensional lattice Λ , $1 \leq \lambda_1(\Lambda) \cdot \lambda_n(\Lambda^*) \leq n$.

36 For $r > 0$, the Gaussian function $\rho_r : H \rightarrow (0, 1]$ is defined to be $\rho_r(\mathbf{x}) = \exp(-\pi \|\mathbf{x}\|^2 / r^2)$. The continuous
 37 Gaussian probability distribution D_r over H is given by the density function $r^{-n} \rho_r(\mathbf{x})$. Note that D_r is the
 38 n -dimensional normal distribution with mean vector $(0, \dots, 0)$ and variance/co-variance matrix $\text{diag}(\sigma^2, \dots, \sigma^2)$
 39 where $\sigma = r/\sqrt{2\pi}$. Consequently, if X_1 and X_2 are independent random variables following D_{r_1} and D_{r_2} , then
 40 $X_1 + X_2$ follows $D_{\sqrt{r_1^2 + r_2^2}}$ (see Theorem 4.2.14 of [9]).

1 For a lattice Λ , a point $\mathbf{u} \in H$ and a positive real r , the discrete Gaussian probability distribution over $\Lambda + \mathbf{u}$
2 with parameter r is defined to be $D_{\Lambda+\mathbf{u},r}(\mathbf{x}) = \rho_r(\mathbf{x})/\rho_r(\Lambda + \mathbf{u})$, where $\rho_r(\Lambda + \mathbf{u})$ denotes $\sum_{\mathbf{y} \in \Lambda+\mathbf{u}} \rho_r(\mathbf{y})$ and
3 more generally $\rho_r(S) = \sum_{\mathbf{y} \in S} \rho_r(\mathbf{y})$ for a countable subset $S \subset H$.

4 For a lattice Λ and a positive real ε , the smoothing parameter $\eta_\varepsilon(\Lambda)$ is defined to be the smallest r such that
5 $\rho_{1/r}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \varepsilon$. Given a lattice $\Lambda = \mathcal{L}(B)$, it can be shown that for any $r \geq \eta_\varepsilon(\Lambda)$, the statistical distance
6 between the uniform distribution on the fundamental parallelepiped $\mathcal{P}(B)$ and the distribution obtained by
7 sampling from D_r and reducing the result modulo the lattice to an element of $\mathcal{P}(B)$ is at most $\varepsilon/2$ (see Lemma 5
8 of [28]). Thus, if one chooses a very small value for ε , then the Gaussian distribution D_r considered over $\mathcal{P}(B)$
9 with $r \geq \eta_\varepsilon(\Lambda)$ behaves essentially like the uniform distribution. As shown in Claim 2.13 of [29], a lower bound
10 for $\eta_\varepsilon(\Lambda)$ in terms of $\lambda_1(\Lambda^*)$ can be obtained by setting ε equal to the term in the sum $\sum_{\mathbf{x} \in \Lambda^* \setminus \{\mathbf{0}\}} \exp(-\pi(\eta\|\mathbf{x}\|)^2)$
11 corresponding to a shortest vector \mathbf{x} in Λ^* . That is, $\eta_\varepsilon(\Lambda) \geq \sqrt{(-\ln \varepsilon)/\pi}/\lambda_1(\Lambda^*)$.

12 **Remark 1.** *Along with spherically symmetric distributions D_r , distributions were considered in [19] in which*
13 *$\|\mathbf{x}\|^2/r^2$ is replaced by $\sum_{i=1}^n a_i^2/r_i^2$, where $\mathbf{x} = a_i \mathbf{h}_i$. In this paper we only use spherically symmetric distributions,*
14 *which suffice for our purposes.*

15 A number field $K = \mathbb{Q}(\zeta)$ is obtained by adjoining ζ to \mathbb{Q} , where ζ is a root of a monic irreducible polynomial
16 $f(x) \in \mathbb{Q}[x]$. The degree of $f(x)$ is the degree of the number field K . Let n denote the degree. Then K is an
17 n -dimensional vector space over \mathbb{Q} . Over \mathbb{C} , $f(x)$ has n roots. Recall that the complex roots come in pairs and
18 let s_1 be the number of real roots and $2s_2$ be the number of complex roots so that $s_1 + 2s_2 = n$. Suppose that the
19 roots are ordered as ζ_1, \dots, ζ_n , where $\zeta_1, \dots, \zeta_{s_1}$ are real and $\zeta_{s_1+s_2+j} = \overline{\zeta_{s_1+j}}$ for $j = 1, \dots, s_2$. Let $\sigma_i : K \rightarrow \mathbb{C}$ be
20 the embedding of K in \mathbb{C} obtained by extending the map $\zeta \mapsto \zeta_i$. The canonical embedding $\sigma : K \rightarrow \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$
21 is given by $\sigma(x) = (\sigma_1(x), \dots, \sigma_n(x))$. Note that for any $x \in K$, and $i = 1, \dots, s_2$, $\sigma_{s_1+s_2+i}(x) = \overline{\sigma_{s_1+i}(x)}$,
22 so that $\sigma(K) \subset H$. For $x \in K$, the trace and norm of x are respectively defined as $\text{Tr}(x) = \sum_{i=1}^n \sigma_i(x)$ and
23 $N(x) = \prod_{i=1}^n \sigma_i(x)$. A geometric norm on an element $x \in K$ is defined via the embedding σ to be $\|x\| = \|\sigma(x)\|$
24 (and $\|x\|_\infty = \|\sigma(x)\|_\infty$).

25 An algebraic integer is a root of a monic polynomial with integer coefficients. Let \mathcal{O}_K denote the set of all
26 algebraic integers in the number field K . Under usual addition and multiplication in K , \mathcal{O}_K forms a ring called
27 the ring of integers of K . Any ideal (also called integral ideal) of \mathcal{O}_K is also a free \mathbb{Z} -module of rank n , i.e., it
28 is generated as the set of all \mathbb{Z} -linear combinations of some basis $\{u_1, \dots, u_n\} \subset \mathcal{O}_K$. The norm of an ideal \mathcal{I}
29 is defined to be $N(\mathcal{I}) = \#(\mathcal{O}_K/\mathcal{I})$. A fractional ideal $\mathcal{I} \subset K$ is a set such that $d\mathcal{I}$ is an integral ideal of \mathcal{O}_K
30 for some $d \in \mathcal{O}_K$. The norm of a fractional ideal \mathcal{I} is defined to be $N(\mathcal{I}) = N(d\mathcal{I})/|N(d)|$. The set of fractional
31 ideals form a group under multiplication.

32 A fractional ideal \mathcal{I} has a \mathbb{Z} -basis $\{u_1, \dots, u_n\}$. Under the canonical embedding $\sigma(\mathcal{I})$ is a lattice, called an
33 ideal lattice having basis $\{\sigma(u_1), \dots, \sigma(u_n)\}$. The fractional ideal \mathcal{I} is identified with its embedding $\sigma(\mathcal{I})$ and
34 one talks of the minimum distance $\lambda_1(\mathcal{I})$ of \mathcal{I} and similarly for other lattice quantities. Likewise, given $r > 0$,
35 $D_{\mathcal{I},r}$ denotes the distribution $D_{\sigma(\mathcal{I}),r}$ over $\sigma(\mathcal{I})$. The (absolute) discriminant Δ_K of K is defined to be the square
36 of the fundamental volume of the ideal lattice $\sigma(\mathcal{O}_K)$. The fundamental volume of any ideal lattice $\sigma(\mathcal{I})$ is
37 $N(\mathcal{I}) \cdot \sqrt{\Delta_K}$.

38 A lattice in K is the \mathbb{Z} -span of a \mathbb{Q} -basis of K . Let Λ be a lattice in K . The conjugate dual of Λ is defined
39 to be $\Lambda^\vee = \{x \in K : \text{Tr}(x\Lambda) \subseteq \mathbb{Z}\}$. It follows that $\sigma(\Lambda^\vee) = \overline{\sigma(\Lambda)^*}$. Let $R = \mathcal{O}_K$ which is a lattice in K . The
40 fractional ideal R^\vee is called the codifferent. For any ideal \mathcal{I} , $\mathcal{I}^\vee = \mathcal{I}^{-1} \cdot R^\vee$.

41
42 **Ideal SVP and SIVP:** An instance of the γ -approximate shortest vector problem for K , denoted $K\text{-SVP}_\gamma$, is
43 a fractional ideal \mathcal{I} in K and it is required to find a nonzero $x \in \mathcal{I}$ such that $\|x\| \leq \gamma \cdot \lambda_1(\mathcal{I})$. The γ -approximate
44 shortest independent vector problem in K , denoted $K\text{-SIVP}_\gamma$, requires finding n linearly independent elements
45 in \mathcal{I} all of whose norms are at most $\gamma \cdot \lambda_n(\mathcal{I})$.

1 **Ring-LWE distribution:** Recall that the intuitive meaning of an LWE sample is an approximate linear equation
2 of the form $\sum a_i s_i = b$ where b includes an error, the s_i are unknown, and the a_i and b are given to the solver. In
3 ring-LWE $\sum a_i s_i$ is realised more efficiently because a and s are elements of a number field, and the left hand side
4 is simply $a \cdot s$. More precisely, let K be a number field and $R = \mathcal{O}_K$. For a fractional ideal \mathcal{J} in K and an integer
5 $q \geq 2$, let \mathcal{J}_q denote the set of residue classes of \mathcal{J} modulo $q\mathcal{J}$. Let $\mathbb{T} = H/\sigma(R^\vee)$ denote H modulo $\sigma(R^\vee)$.
6 Suppose $s \in R_q^\vee$ and $a \in R_q$. There are elements $x \in R^\vee$ and $y \in R$ such that $s = x + qR^\vee$ and $a = y + qR$.
7 Define the result of the operation $a \cdot s$ to be $xy + qR^\vee$ which is in R_q^\vee . One can show that the operation is well
8 defined. Similarly, the result of the operation $(a \cdot s)/q$ is defined to be $xy/q + R^\vee$ which is in $(1/q)R^\vee$ modulo
9 R^\vee . By $\sigma((a \cdot s)/q)$ we will denote the element $\sigma(xy/q) + \sigma(R^\vee)$ of \mathbb{T} . For $s \in R_q^\vee$ and a positive real number r ,
10 a sample from the ring-LWE distribution $A_{s,r}$ over $R_q \times \mathbb{T}$ is $(a, \sigma((a \cdot s)/q) + \mathbf{e} \bmod \sigma(R^\vee))$, where a is chosen
11 uniformly at random from R_q and \mathbf{e} is chosen from H following the distribution D_r .

12 **Remark 2.** 1. We have defined the ring-LWE distribution by transferring $(a \cdot s)/q$ to H and performing the
13 addition with the error \mathbf{e} in H modulo $\sigma(R^\vee)$. This is helpful for the theoretical analysis of the reduction. In
14 practice, on the other hand, it is computationally more efficient to transfer \mathbf{e} to an approximate element in
15 K and perform the addition in K . We briefly describe how $\mathbf{e} \in H$ can be transferred to K . Let (ν_1, \dots, ν_n)
16 be a basis of K over \mathbb{Q} . The requirement is to find $x_1, \dots, x_n \in \mathbb{Q}$, such that $\sigma(\sum_{i=1}^n x_i \nu_i) = \sum_{i=1}^n x_i \sigma(\nu_i)$
17 is close to \mathbf{e} . This is done as follows. Let M be the inverse of the matrix whose (i, j) -th entry is $\sigma_j(\nu_i)$,
18 compute $(y_1, \dots, y_n) = \mathbf{e}M \in \mathbb{R}^n$ and then choose x_i to be a rational approximation of y_i .

19 2. In [19], the second component of the ring-LWE distribution is an element of the field tensor product $K \otimes_{\mathbb{Q}} \mathbb{R}$.
20 Since $K \otimes_{\mathbb{Q}} \mathbb{R}$ and H are isomorphic as n -dimensional vector spaces over \mathbb{R} , we have chosen to work with
21 the equivalent and simpler formulation where the second component of the ring-LWE is an element of H .

22 **Search ring-LWE:** Let $\alpha > 0$ be a real number and $q \geq 2$ be an integer. The search version of the ring-LWE
23 problem, denoted ring-LWE $_{q, \leq \alpha}$, is the following. For any $s \in R_q^\vee$ and a fixed positive real number $r \leq \alpha$, given
24 access to arbitrarily many independent samples from $A_{s,r}$, find s . Formally, a probabilistic algorithm \mathcal{A} to solve
25 ring-LWE $_{q, \leq \alpha}$ has access to an oracle $W_{s,r}$, where $r \leq \alpha$ is unknown, which when queried returns an independent
26 sample from $A_{s,r}$. \mathcal{A} is allowed to adaptively query $W_{s,r}$ a number of times and at the end outputs an element
27 $s' \in R_q^\vee$. The success probability of \mathcal{A} is the probability that $s' = s$. The important parameters for \mathcal{A} are its
28 success probability, its runtime and the number of times it queries its oracle.

29 A necessary condition for solvability of ring-LWE $_{q, \leq \alpha}$ is $\alpha < \eta_\varepsilon(R^\vee)$ for all negligible ε , as otherwise the added
30 error makes the samples essentially uniform.

31
32 **Ring DLWE (fixed width):** Let $r > 0$ be a real number and $q \geq 2$ be an integer. The decision version of
33 the ring-LWE problem, denoted ring-DLWE $_{q,r}$, is the following. Let s be chosen uniformly at random from R_q^\vee .
34 The task is to distinguish with non-negligible advantage between arbitrarily many independent samples from
35 $A_{s,r}$ and the same number of samples generated independently and uniformly from $R_q \times \mathbb{T}$. Formally, let \mathcal{D} be a
36 distinguisher which takes as input a list \mathcal{T} consisting of elements from $R_q \times \mathbb{T}$. For a fixed value of $s \in R_q^\vee$, let $p_{s,0}$
37 be the probability that \mathcal{D} outputs 1 when \mathcal{T} consists of independent samples from $A_{s,r}$, where the probability is
38 over all components of the input other than s as well as the internal coin tosses of \mathcal{D} . Let p_1 be the probability
39 that \mathcal{D} outputs 1 when \mathcal{T} consists of independent samples chosen uniformly from $R_q \times \mathbb{T}$, where the probability
40 is over the input and the internal coin tosses of \mathcal{D} . For a given value of s , the advantage of the distinguisher is
41 $|p_{s,0} - p_1|$. For $\epsilon_1, \epsilon_2 \in (0, 1]$, we say that \mathcal{D} is an (ϵ_1, ϵ_2) -distinguisher if \mathcal{D} has advantage at least ϵ_2 for at least
42 a proportion ϵ_1 of the set of possible $s \in R_q^\vee$,

43 **Remark 3.** A more general definition of ring DLWE was given in [19]. In this definition, the error distribution
44 itself is chosen from a distribution over a family of error distributions. Theorem 5.1 of [19] was proved for
45 such a definition, while Theorem 5.2 of [19] was proved for the case where the error distribution is fixed. Since
46 applications use a fixed error distribution, we have defined ring DLWE with fixed error distribution.

1 **Notation** We summarise and fix some notation for future convenience.

$\log(x)$: logarithm of x to the base 2
$\ln(x)$: natural logarithm of x
Λ	: a lattice in H
$\lambda_1(\Lambda)$: minimum distance of the lattice Λ
$\lambda_n(\Lambda)$: the least real number such that Λ has n linearly independent vectors with the length of the longest being equal to this number
$\eta_\varepsilon(\Lambda)$: smoothing parameter for a lattice Λ
K	: underlying number field
n	: degree of the number field
\mathcal{O}_K, R	: ring of integers of K
\mathcal{I}, \mathcal{J}	: fractional ideals of a number field
σ	: canonical embedding of a number field into H (as defined in (3))
\mathbb{T}	: the set of residue classes of H modulo $\sigma(R^\vee)$
\mathcal{I}^*	: the dual of the lattice $\sigma(\mathcal{I})$
\mathcal{I}^\vee	: the conjugate dual of the lattice $\sigma(\mathcal{I})$
D_r	: Gaussian distribution on H of width r
$D_{\mathcal{I},r}$: discrete Gaussian distribution of width r on the lattice $\sigma(\mathcal{I})$
$r, r', r_i, \xi, \xi', \xi_i$: widths of Gaussian distributions
q	: an integer ≥ 2
\mathcal{I}_q	: the set of residue classes of \mathcal{I} modulo $q\mathcal{I}$
$\omega(\sqrt{\ln n})$: a fixed function which grows asymptotically faster than $\sqrt{\ln n}$
α	: an upper bound on the width of Gaussian distributions that satisfies $\alpha < \sqrt{\ln n/n}$ and $\alpha q \geq 2\omega(\sqrt{\ln n})$.

2 3 Reducing K -SIVP $_\gamma$ to search ring-LWE $_{q,\leq\alpha}$

3 Fix three parameters: a positive integer n which will denote the degree of the number field K ; an integer $q \geq 2$
4 which is used to define the ring-LWE problem; and a positive real number α such that $\alpha q \geq 2 \cdot \omega(\sqrt{\ln n})$. We
5 assume $\alpha < \sqrt{\ln n/n}$ which, as remarked in [19], holds for proposed applications. In the asymptotic setting, q
6 and α are considered to be functions of n .

7 The reduction of K -SIVP $_\gamma$ to ring-LWE $_{q,\leq\alpha}$ is obtained by composing reductions involving several intermediate
8 computational problems. Let K be a number field.

- 9 • Let Γ be a function from fractional ideals in K to \mathbb{R} . The discrete Gaussian sampling problem in K ,
10 denoted K -DGS $_\Gamma$, is the following. Given a fractional ideal \mathcal{I} in K and $r \geq \Gamma(\mathcal{I})$, produce a sample from
11 $D_{\mathcal{I},r}$. This means producing an element of \mathcal{I} such that the probability of producing $x \in \mathcal{I}$ is given by
12 $D_{\mathcal{I},r}(x) = \rho_r(\sigma(x)) / \sum_{y \in \mathcal{I}} \rho_r(\sigma(y))$.
- 13 • Given a fractional ideal \mathcal{I} in K and $\xi < \lambda_1(\mathcal{I}) / (2\sqrt{2n})$, an instance of the bounded distance decoding
14 problem K -BDD $_{\mathcal{I},\xi}$ is an element $y \in K$ such that $y = x + e$, where $x \in \mathcal{I}$ and $e = \sigma^{-1}(\mathbf{e})$ is chosen
15 according to D_ξ ; the requirement is to find x' such that $x' = x$ except with negligible probability where the
16 probability is over \mathbf{e} as well as internal coin tosses. An algorithm to solve K -BDD $_{\mathcal{I},\xi}$ will take as input the
17 pair (\mathcal{I}, y) . The upper bound on ξ ensures that the solution x is unique except with negligible probability.
18 Note that ξ is unknown to the solver.
- 19 • For a fractional ideal \mathcal{I} in K and an integer $q \geq 2$, the q -BDD $_{\mathcal{I},\xi}$ problem is the following. Given an instance
20 y of K -BDD $_{\mathcal{I},\xi}$ with solution $x \in \mathcal{I}$, find $x \bmod q\mathcal{I}$.

1 The K -SIVP $_\gamma$ to ring-LWE $_{q,\leq\alpha}$ reduction is obtained from the following sequence of algorithms, in which \mathcal{A}_i calls
 2 \mathcal{A}_{i+1} as an oracle, for $0 \leq i \leq 4$.

3
 4 **Algorithm \mathcal{A}_0 :** Solves K -SIVP $_\gamma$ for an appropriate value of γ (see §3.2 below). The input is a fractional ideal
 5 \mathcal{I} and the output is a set of n linearly independent elements of \mathcal{I} the longest of which is at most $\gamma\lambda_n(\mathcal{I})$.

6
 7 **Algorithm \mathcal{A}_1 :** Solves K -DGS $_\Gamma$, for an appropriate Γ (see §3.2 below). The input is a pair (\mathcal{I}, r) , where \mathcal{I} is a
 8 fractional ideal of K and $r \geq \Gamma(\mathcal{I})$. The output is a sample from the distribution $D_{\mathcal{I},r}$.

9
 10 **Algorithm \mathcal{A}_2 :** This is a quantum algorithm which, given as input a fractional ideal \mathcal{I} and a set of samples
 11 chosen independently from $D_{\mathcal{I},r}$, returns a sample from $D_{\mathcal{I},r'}$, where $r' \leq r/2$. The conditions on r and r' are
 12 given in (27) of Appendix A.

13
 14 **Algorithm \mathcal{A}_3 :** Solves K -BDD $_{\mathcal{I}^\vee,\xi}$. The input is a pair (\mathcal{I}^\vee, y) , where \mathcal{I} is a fractional ideal of K , $y = x + e$,
 15 $x \in \mathcal{I}^\vee$, and $e = \sigma^{-1}(\mathbf{e})$ is chosen according to the distribution D_ξ . Additionally, \mathcal{A}_3 has access to a set of
 16 samples chosen independently from $D_{\mathcal{I},r}$. The output is an $x' \in \mathcal{I}^\vee$ such that $x' = x$ except with negligible
 17 probability. The relation between r and ξ is given in (27) of Appendix A.

18
 19 **Algorithm \mathcal{A}_4 :** Solves q -BDD $_{\mathcal{I}^\vee,\xi}$. The input to \mathcal{A}_4 is the same as that to \mathcal{A}_3 and the output is $x' \bmod q$ such
 20 that $x' \equiv x \bmod q$ except with negligible probability.

21
 22 **Algorithm \mathcal{A}_5 :** Solves ring-LWE $_{q,\leq\alpha}$. Algorithm \mathcal{A}_5 has access to an oracle which generates samples from the
 23 ring-LWE distribution $A_{s,r}$ defined in §2, where $r \leq \alpha$ and both s and r are unknown to \mathcal{A}_5 . The algorithm
 24 interacts with the oracle and finally outputs s .

25 3.1 Reducing K -SIVP $_\gamma$ to K -DGS $_\Gamma$

26 For a fractional ideal \mathcal{I} in K , let

$$\Gamma(\mathcal{I}) = \frac{\gamma \cdot \lambda_n(\mathcal{I})}{2\sqrt{n}}. \quad (4)$$

27 Suppose $\varepsilon \leq 1/10$ and $\gamma \geq 2\sqrt{2n}\eta_\varepsilon(\mathcal{I})/\lambda_n(\mathcal{I})$, which implies that $\Gamma(\mathcal{I}) \geq \sqrt{2}\eta_\varepsilon(\mathcal{I})$. Given an algorithm \mathcal{A}_1 to
 28 solve K -DGS $_\Gamma$ where Γ is given by (4), it is possible to construct an algorithm \mathcal{A}_0 to solve K -SIVP $_\gamma$. This is
 29 shown in Lemma 3.17 of [29].

30 We briefly review the proof. The objective is to obtain a set of n linearly independent vectors whose longest
 31 vector has length at most $2\sqrt{n}\Gamma(\mathcal{I})$, which using (4) is equal to $\gamma\lambda_n(\mathcal{I})$. Algorithm \mathcal{A}_0 uses the LLL algorithm
 32 to obtain a set B_0 of n linearly independent vectors such that the length d_0 of the longest vector in B_0 satisfies
 33 $d_0 \leq 2^{(n-1)/2}\lambda_n(\mathcal{I})$ (see Theorem 3 of [21]). From $d_0 \leq 2^{(n-1)/2}\lambda_n(\mathcal{I}) < 2^{n/2}\lambda_n(\mathcal{I})$ and (4), we have $d_0 < 2^n\Gamma(\mathcal{I})$
 34 for $\gamma \geq \sqrt{n}/2^{n/2-1}$, where the condition on γ holds⁴ for all $n \geq 8$.

35 For $i = 1, \dots, n$, let $u_i = d_0/2^{i-1}$. For each i in $\{1, \dots, n\}$, \mathcal{A}_0 does the following. It invokes \mathcal{A}_1 a total of
 36 n^2 times on the input (\mathcal{I}, u_i) to obtain a set T_i of n^2 elements of \mathcal{I} chosen independently from the distribution⁵
 37 $D_{\mathcal{I},u_i}$. \mathcal{A}_0 looks for a set B_i of n linearly independent elements in T_i . If B_i is found, then let d_i be the length of

⁴From (4), the assumptions $\Gamma(\mathcal{I}) \geq \sqrt{2}\eta_\varepsilon(\mathcal{I})$ and $\varepsilon \leq 1/10$ and the lower bound on $\eta_\varepsilon(\mathcal{I})$ given in Claim 2.13 of [29], it follows
 that $\gamma \geq \sqrt{8 \ln 10 / (n\pi)} \geq \sqrt{n}/2^{n/2-1}$ for $n \geq 8$.

⁵ If $d_0 < \Gamma(\mathcal{I})$, let $j = 0$, otherwise, let $j \in \{1, \dots, n\}$ be such that $2^{j-1}\Gamma(\mathcal{I}) \leq d_0 < 2^j\Gamma(\mathcal{I})$. Then for $j+1 \leq i \leq n$, we have
 $u_i < \Gamma(\mathcal{I})$ implying that the pair (\mathcal{I}, u_i) is an invalid input to \mathcal{A}_1 . Since the expression for $\Gamma(\mathcal{I})$ given by (4) involves $\lambda_n(\mathcal{I})$, the
 value of $\Gamma(\mathcal{I})$ cannot be efficiently computed and so it is not possible to check the condition $u_i < \Gamma(\mathcal{I})$. We redefine \mathcal{A}_1 so that if
 $u_i < \Gamma(\mathcal{I})$, it returns some element of \mathcal{I} , but with nothing assumed about whether the selection adheres to any Gaussian distribution.

1 the longest vector in B_i . Finally \mathcal{A}_0 returns a set B_k such that d_k is the minimum of all the d_i such that B_i is
 2 defined.

3 The claim is that with high probability \mathcal{A}_0 returns a set of n linearly independent vectors whose longest
 4 vector is at most $2\sqrt{n}\Gamma(\mathcal{I})$. Since $d_0 < 2^n\Gamma(\mathcal{I})$, it follows that either $d_0 < \Gamma(\mathcal{I})$, or there is some $k_0 \in \{1, \dots, n\}$
 5 such that $\Gamma(\mathcal{I}) \leq u_{k_0} < 2\Gamma(\mathcal{I})$. If $d_0 < \Gamma(\mathcal{I})$, then $d_k \leq d_0 < \Gamma(\mathcal{I})$ and the claim holds. Otherwise, consider the
 6 k_0 such that $\Gamma(\mathcal{I}) \leq u_{k_0} < 2\Gamma(\mathcal{I})$. The conditions $\varepsilon \leq 1/10$ and $u_{k_0} \geq \Gamma(\mathcal{I}) \geq \sqrt{2}\eta_\varepsilon(\mathcal{I})$ ensure that with high
 7 probability the set T_{k_0} of n^2 vectors contains a set B_{k_0} of n linearly independent vectors (Corollary 3.16 of [29]).
 8 Further, with high probability the vectors in B_{k_0} are of length at most $u_{k_0}\sqrt{n}$ (Lemma 2.5 of [29]) which is less
 9 than $2\sqrt{n}\Gamma(\mathcal{I})$, i.e., $d_{k_0} < 2\sqrt{n}\Gamma(\mathcal{I})$. Since $d_k \leq d_{k_0}$, the claim also holds in this case. We record the following.

10 **Proposition 1.** \mathcal{A}_0 invokes \mathcal{A}_1 a total of n^3 times.

11 3.2 Reducing K -DGS $_\Gamma$ to ring-LWE $_{q,\leq\alpha}$

12 For a fractional ideal \mathcal{I} in K , let

$$\gamma = 2\sqrt{2} \cdot \frac{\sqrt{n}}{\alpha} \cdot \omega(\sqrt{\ln n}) \cdot \frac{\eta_\varepsilon(\mathcal{I})}{\lambda_n(\mathcal{I})} \quad (5)$$

13 Since $\alpha < \sqrt{\ln n/n}$ and $\omega(\sqrt{\ln n})$ grows faster than $\sqrt{\ln n}$, the value of γ given by (5) satisfies the lower bound
 14 on γ assumed in Section 3.1. Substituting the value of γ given by (5) in (4), we obtain

$$\Gamma(\mathcal{I}) = \frac{\sqrt{2} \cdot \omega(\sqrt{\ln n}) \cdot \eta_\varepsilon(\mathcal{I})}{\alpha}. \quad (6)$$

15 With Γ given by (6), there is a reduction from K -DGS $_\Gamma$ to ring-LWE $_{q,\leq\alpha}$, where $\alpha q \geq 2\omega(\sqrt{\ln n})$. The reduction
 16 will be described in this section.

17 For γ given by (5), using Lemma 2.2⁶ of [19] we have

$$\gamma \leq 2\sqrt{2} \cdot \frac{\sqrt{n}}{\alpha} \cdot \omega(\sqrt{\ln n}) \sqrt{\ln(n/\varepsilon)}. \quad (7)$$

18 The relation (7) has been summarised in [19] as $\gamma = \tilde{O}(\sqrt{n}/\alpha)$.

19 **Remark 4.** As ε decreases, the expression in (7) increases. If $\varepsilon = e^{-n}$, the order of magnitude of γ is $\tilde{O}(n/\alpha)$
 20 and not $\tilde{O}(\sqrt{n}/\alpha)$. The expression $\gamma = \tilde{O}(\sqrt{n}/\alpha)$ in [19] implies that ε is assumed not to be too small, although
 21 it has to be a negligible function of n . More precisely, suppose there is some constant d such that

$$\varepsilon \geq n^{1-(\log n)^d}. \quad (8)$$

22 Then $\gamma = \tilde{O}(\sqrt{n}/\alpha)$ is justified.

23 We now consider the reduction of K -DGS $_\Gamma$ to ring-LWE $_{q,\leq\alpha}$ in further detail. While going through the
 24 description below, it will be helpful to keep in mind the inputs and outputs of the various algorithms which are
 25 shown in Figure 1.

26 The input to \mathcal{A}_1 is a pair (\mathcal{I}, r) , where \mathcal{I} is an ideal and $r \geq \Gamma(\mathcal{I})$. Let $i_0 = 2n + \lceil (\log n)/2 \rceil$. For $i = 0, \dots, i_0$,
 27 define $r_i = r \cdot (\alpha q / \omega(\sqrt{\ln n}))^i$. Note that $r_i \geq 2^i r$ because, by the above assumption, $\alpha q \geq 2\omega(\sqrt{\ln n})$. Also,
 28 $r_{i_0} \geq 2^{2n} \sqrt{n} r$. The lower bound on r_{i_0} ensures that \mathcal{A}_1 can sample from $D_{\mathcal{I}, r_{i_0}}$ without requiring the help of the
 29 LWE oracle (see Lemma 3.2 of [29]). For $i = 1, \dots, i_0$, define $\xi_i = (\alpha q) / (2r_i \cdot \omega(\sqrt{\ln n}))$.

⁶Lemma 2.2 of [19] shows that $\eta_\varepsilon(\mathcal{I}) \leq \sqrt{\ln(n/\varepsilon)} \lambda_n(\mathcal{I})$.

\mathcal{A}_1 solves $K\text{-DGS}_\Gamma$ <ul style="list-style-type: none"> • input (\mathcal{I}, r), $r \geq \Gamma(\mathcal{I})$ • output a “sample” from $D_{\mathcal{I},r}$ 	\mathcal{A}_2 <ul style="list-style-type: none"> • input \mathcal{S} from $D_{\mathcal{I},r}$ • output a “sample” from $D_{\mathcal{I},r'}$ • $r' = r \cdot \omega(\sqrt{\ln n})/(\alpha q) < r/2$.
\mathcal{A}_3 solves $K\text{-BDD}_{\mathcal{I}^\vee, \xi}$ <ul style="list-style-type: none"> • input y such that $y = x + e$, $x \in \mathcal{I}^\vee$, $e = \sigma^{-1}(\mathbf{e})$, \mathbf{e} from D_ξ • has access to \mathcal{S} from $D_{\mathcal{I},r}$ • output $x' \in \mathcal{I}^\vee$ where $x' = x$ almost always • $\xi = \alpha q/(2r\omega(\sqrt{\ln n}))$ 	\mathcal{A}_4 solves $q\text{-BDD}_{\mathcal{I}^\vee, \xi}$ <ul style="list-style-type: none"> • input same as \mathcal{A}_3 • outputs $x' \bmod q$ <hr/> \mathcal{A}_5 solves ring-LWE $_{q, \leq \alpha}$ <ul style="list-style-type: none"> • has access to $A_{s,r}$-oracle • $r \leq \alpha$ and s are unknown to \mathcal{A}_5 • output s

Figure 1: Inputs and outputs of algorithms \mathcal{A}_1 to \mathcal{A}_5 .

1 First, \mathcal{A}_1 prepares a list of N samples \mathcal{S}_{i_0} from $D_{\mathcal{I},r_{i_0}}$. Then for each i starting from i_0 down to 1, \mathcal{A}_1 invokes
2 \mathcal{A}_2 a total of N times, providing it with \mathcal{S}_i , where each call to \mathcal{A}_2 returns a sample from $D_{\mathcal{I},r_{i-1}}$. To obtain a
3 sample from $D_{\mathcal{I},r_{i-1}}$, \mathcal{A}_2 creates a quantum circuit and a quantum state that will produce the desired sample
4 provided it can “uncompute” a nearest vector that is in the first register.

5 To accomplish the erasure of the first entangled register, we need a circuit of gates that reverse the gates in
6 the circuit for \mathcal{A}_3 . All of the algorithms after \mathcal{A}_2 will have to be incorporated into a quantum circuit that must
7 then be reversed so that the resulting quantum circuit can be included in the quantum part of \mathcal{A}_2 . This raises
8 feasibility issues that we will discuss in §6.

9 In \mathcal{A}_3 the offset for the BDD instance is sampled from the distribution D_{ξ_i} . \mathcal{A}_3 solves the BDD instance (via
10 \mathcal{A}_4) by invoking \mathcal{A}_5 . For this, \mathcal{A}_4 needs to be able to simulate the responses to the LWE queries made by \mathcal{A}_5 .

11 The ℓ_∞ distance of a sample from D_{ξ_i} has length at most $\xi'_i = \alpha q/(\sqrt{2}r_i)$ except with negligible probability
12 and this is required in Lemma 4.7 of [19] to show that the simulation of responses to the LWE queries by \mathcal{A}_4 is
13 correct. Below we provide further details of algorithms \mathcal{A}_3 and \mathcal{A}_4 .

14 Algorithm \mathcal{A}_3 takes as input an ideal \mathcal{I}^\vee and an element $y \in K$ such that $y = x + e$, where $x \in \mathcal{I}^\vee$ and
15 $e = \sigma^{-1}(\mathbf{e})$ is chosen according to D_ξ . Additionally, it has access to a set of independent samples from $D_{\mathcal{I},r}$.
16 Algorithm \mathcal{A}_3 returns x . We provide a brief overview of the construction of \mathcal{A}_3 using \mathcal{A}_4 as an oracle. Let B be
17 a matrix whose columns form a basis for the lattice $\sigma(\mathcal{I}^\vee)$. Recall that the only difference between \mathcal{A}_3 and \mathcal{A}_4
18 is that the latter only finds the nearest lattice vector modulo q .

19 We first apply $\mathcal{A}_4(\mathcal{I}^\vee, \mathcal{S}, y)$ to find an integer vector b_1 such that $b_1 \equiv a_1 \pmod q$, where a_1 (which we do not
20 know yet) is the integer vector such that the lattice element Ba_1 is the closest vector in $\sigma(\mathcal{I}^\vee)$ to σy . Set $y_1 = y$.
21 Now repeat the procedure with $y_2 = (y_1 - Bb_1)/q$, obtaining $b_2 \equiv a_2 \pmod q$ such that Ba_2 is the closest vector
22 in $\sigma(\mathcal{I}^\vee)$ to $\sigma(y_2)$. Since $B(a_1 - b_1)$ is the closest vector to $y_1 - Bb_1 = qy_2$, it follows that $a_2 = (a_1 - b_1)/q$, and
23 that the distance between Ba_2 and y_2 is less than the distance between a_1 and b_1 by a factor of q . Applying \mathcal{A}_4
24 $n - 1$ times – generating the sequences y_i and b_i , $i = 1, \dots, n - 1$ – we finally get y_n close enough to the lattice
25 that we can use Babai’s nearest plane algorithm to find the nearest lattice vector to y_n , which is a_n . Once we
26 know a_n , as well as the b_1, \dots, b_n , we can successively compute $a_i = b_i + qa_{i+1}$, $i = n - 1, \dots, 1$. Then Ba_1 is
27 the desired output of \mathcal{A}_3 .

28 The input to algorithm \mathcal{A}_4 is the same as that to algorithm \mathcal{A}_3 . \mathcal{A}_4 returns $x \bmod q\mathcal{I}$. The general task of
29 \mathcal{A}_4 in solving an instance of $q\text{-BDD}_{\mathcal{I}^\vee, \xi}$ is similar to that of \mathcal{A}_5 in solving ring-LWE $_{q, \leq \alpha}$. But there are two major
30 differences. First, \mathcal{A}_4 works with the lattice \mathcal{I}^\vee modulo q , whereas \mathcal{A}_5 works with the lattice R^\vee modulo q . The
31 reduction handles this by using an element $t \in \mathcal{I}$ that gives an isomorphism from $R \bmod q$ to $\mathcal{I} \bmod q$ and also
32 in the other direction between the dual lattices $\mathcal{I}^\vee \bmod q$ and $R^\vee \bmod q$. The second difference is that \mathcal{A}_4 gets
33 just one input vector $y = x + e$, whereas \mathcal{A}_5 has access to an oracle that provides N samples from $A_{s,r}$. Since

1 \mathcal{A}_4 calls \mathcal{A}_5 , the oracle queries made by \mathcal{A}_5 has to be simulated by \mathcal{A}_4 . This is done by randomising the error in
 2 y , that is, by adding errors \mathbf{e}' chosen according to $D_{\alpha/\sqrt{2}}$. For details, see §4.2 of [19].

3 Based on the overview and the above descriptions of algorithms $\mathcal{A}_1, \mathcal{A}_3$ and \mathcal{A}_4 , we record the following.

- 4 **Proposition 2.** 1. \mathcal{A}_1 invokes \mathcal{A}_2 a total of $(2n + \lceil (\log n)/2 \rceil)N$ times, where N is the number of LWE
 5 samples required by \mathcal{A}_5 .
- 6 2. \mathcal{A}_2 invokes the reverse circuit of \mathcal{A}_3 once.
- 7 3. \mathcal{A}_3 invokes \mathcal{A}_4 a total of n times.
- 8 4. \mathcal{A}_4 invokes \mathcal{A}_5 once.

9 The reduction of $K\text{-DGS}_\Gamma$ to ring-LWE $_{q,\leq\alpha}$ in [19] is based on the reduction of the DGS problem to the search
 10 LWE problem in [29]. There are, however, some important differences. A number of these differences pertain to
 11 the algebraic techniques needed to handle ideal lattices in [19] that do not apply to the general lattices considered
 12 in [29].

13 One such difference is in the distribution of the error in an LWE sample. In [29] the error follows a fixed width
 14 Gaussian distribution, while in [19] the error follows a distribution drawn from a family of elliptical Gaussian
 15 distributions. Elliptical, rather than spherical, distributions are needed in [19] to argue for the correctness of the
 16 distribution of the error in the simulated LWE samples arising in the $q\text{-BDD}_{\mathcal{I}^\vee, \xi}$ to ring-LWE $_{q,\leq\alpha}$ reduction. It
 17 appears that the use of elliptical distributions is a mathematical artifact introduced for the sake of the reduction
 18 rather than being of any practical importance.

19 3.3 The tightness gap in the $K\text{-SIVP}_\gamma$ to ring-LWE $_{q,\leq\alpha}$ reduction

20 The following theorem is a concrete version of the reduction of $K\text{-SIVP}_\gamma$ to ring-LWE $_{q,\leq\alpha}$ that is described
 21 in Theorem 4.1 of [19] and in the discussion following that theorem. The tightness gap in (9) follows from
 22 Propositions 1 and 2 above.

23 **Theorem 3.** Let K be an arbitrary number field of degree n , $q \geq 2$ be a positive integer and α be a positive real
 24 number such that $\alpha q \geq 2\omega(\sqrt{\ln n})$ and $\alpha < \sqrt{\ln n/n}$. There is a quantum reduction using approximately $3n^2$
 25 logical qubits from $K\text{-DGS}_\Gamma$, where Γ is given by (6) with $\varepsilon \leq e^{-\pi}$, to ring-LWE $_{q,\leq\alpha}$. Additionally, suppose there
 26 is a positive constant d , such that $\varepsilon \geq n^{1-(\log n)^d}$. Then there is a quantum reduction from $K\text{-SIVP}_{\tilde{O}(\sqrt{n}/\alpha)}$ to
 27 ring-LWE $_{q,\leq\alpha}$.

28 An algorithm \mathcal{A}_0 to solve $K\text{-SIVP}_{\tilde{O}(\sqrt{n}/\alpha)}$ can be constructed using an algorithm \mathcal{A}_5 to solve ring-LWE $_{q,\leq\alpha}$
 29 and the number of times \mathcal{A}_0 calls \mathcal{A}_5 is approximately

$$(2n + (\log n)/2)n^4 \cdot N, \tag{9}$$

30 where N is the number of ring-LWE samples required by \mathcal{A}_5 .

31 **Remark 5.** In arriving at the expression in (9), we have assumed that the time taken by a reverse circuit for
 32 \mathcal{A}_3 is the same as the time taken by a circuit for \mathcal{A}_3 .

33 The term N in (9) is the number of samples that is required in ring-LWE. In practice, N would depend on
 34 α . If α is very small, then the error in the ring-LWE distribution is also very small leading to a relatively easy
 35 instance of ring-LWE. For an easy instance, obtaining about n samples might be sufficient to solve the ring-LWE
 36 problem. With a larger α , the error in the ring-LWE distribution would also be larger, which suggests that the
 37 number of samples required to solve the ring-LWE problem would be larger. Consequently, the number of oracle
 38 queries (and hence the tightness gap) grows with the difficulty of the ring-LWE problem. More concretely, if we

1 assume $N = n^c$, then c increases as α increases. If α is small enough so that ring-LWE is easy and the tightness
2 gap is less, then $K\text{-SIVP}_\gamma$ has a larger approximation factor γ , and so becomes easier. On the other hand, as α
3 increases, causing ring-LWE $_{q,\leq\alpha}$ intuitively to become harder (the case that is relevant for cryptography), then
4 the $K\text{-SIVP}_\gamma$ problem that reduces to it becomes harder (since γ decreases), but its connection to ring-LWE $_{q,\leq\alpha}$
5 becomes weaker because of the larger tightness gap as c increases. Because of this loosening of the connection
6 between the two problems, it would be consistent with the reduction for the ring-LWE $_{q,\leq\alpha}$ search problem to
7 become harder as α increases but at a slower rate than the $K\text{-SIVP}_{\tilde{O}(\sqrt{n}/\alpha)}$ problem. This once again shows that
8 what the reduction gives us is somewhat less than it might seem at first glance.

9 4 Reducing search ring-LWE to ring-DLWE

10 The reduction makes use of several algebraic properties of number fields that are satisfied, in particular, by
11 cyclotomic number fields. Before getting into the reduction, we briefly mention the relevant algebraic properties
12 of cyclotomic number fields.

13 Let K be a number field and τ be an automorphism of K . One may consider τ to act on $\sigma(K)$ as follows:
14 for $a \in K$, $\tau(\sigma(a)) = \sigma(\tau(a))$. It is possible to extend the action of τ to the whole of H in the following manner.
15 Since τ is the identity map on \mathbb{Q} , it follows that τ is a linear transformation of K (considered as a vector space
16 over \mathbb{Q}) to itself. If we fix a \mathbb{Q} -basis $\{\nu_1, \dots, \nu_n\}$ of K , then the action of τ is given by an $n \times n$ non-singular
17 matrix T whose entries are from \mathbb{Q} . More specifically, for $a \in K$, suppose $a = a_1\nu_1 + \dots + a_n\nu_n$, with $a_i \in \mathbb{Q}$
18 for $i = 1, \dots, n$. Then $\tau(a) = b_1\nu_1 + \dots + b_n\nu_n$, where $(b_1, \dots, b_n)^\top = T(a_1, \dots, a_n)^\top$. Extension of τ to H is
19 done using the matrix T . Note that $\{\sigma(\nu_1), \dots, \sigma(\nu_n)\}$ is an \mathbb{R} -basis of H so that any $\mathbf{x} \in H$ can be written as
20 $\mathbf{x} = x_1\sigma(\nu_1) + \dots + x_n\sigma(\nu_n)$, where $x_i \in \mathbb{R}$ for $i = 1, \dots, n$. Then $\tau(\mathbf{x})$ is defined to be $\mathbf{y} = y_1\sigma(\nu_1) + \dots + y_n\sigma(\nu_n)$,
21 where $(y_1, \dots, y_n)^\top = T(x_1, \dots, x_n)^\top$. Extending an automorphism τ of K to H allows us to apply τ to samples
22 drawn from an error distribution defined over H .

23
24 **Cyclotomic number fields.** For $m \geq 1$, let $\Phi_m(x)$ be the m -th cyclotomic polynomial having degree $n = \varphi(m)$.
25 The m -th cyclotomic number field K is $\mathbb{Q}(\zeta)$, where ζ is a root of $\Phi_m(x)$. Let $R = \mathcal{O}_K$. Henceforth, only
26 cyclotomic number fields will be considered.

27 The power basis $\{1, \zeta, \zeta^2, \dots, \zeta^{n-1}\}$ for K over \mathbb{Q} is also a \mathbb{Z} -basis for $\mathcal{O}_K = \mathbb{Z}[x]/\Phi(x)$. Let \mathcal{I} be a fractional
28 ideal in \mathcal{O}_K and v be a shortest nonzero element in \mathcal{I} . Multiplying v by $1, \zeta, \dots, \zeta^{n-1}$ gives a set of n linearly
29 independent vectors of the same length, and hence $\lambda_n(\mathcal{I}) = \lambda_1(\mathcal{I})$. Consequently, a solution to $K\text{-SIVP}_\gamma$ provides
30 a solution to $K\text{-SVP}_\gamma$ and vice versa. This may be contrasted with the case for general lattices, where it has
31 been shown in [20] that $\text{SIVP}_{\sqrt{n}\gamma}$ reduces in polynomial time to SVP_γ , but not that SIVP_γ reduces to SVP_γ .

32 Let q be a prime number such that $q = 1 \pmod m$ so that $q = km + 1$ for some non-negative integer k . Noting
33 that $\mathbb{Z}_q^* = \langle g \rangle$ for a generator g , it follows that the element $\omega = g^k$ has order m in \mathbb{Z}_q^* . The m -th cyclotomic
34 polynomial factors over \mathbb{Z}_q as $\Phi_m(x) = \prod_{i \in \mathbb{Z}_m^*} (x - \omega^i)$. Consequently, $\langle q \rangle = \prod_{i \in \mathbb{Z}_m^*} \mathfrak{q}_i$, where $\mathfrak{q}_i = \langle q, x - \omega^i \rangle$ is
35 a prime ideal of \mathcal{O}_K having norm q . (Note that the ideals \mathfrak{q}_i have been indexed by elements of \mathbb{Z}_m^* rather than
36 by the integers $\{1, \dots, n\}$.)

37 The field K has n automorphisms $\tau_k(\zeta) = \zeta^k$, for $k \in \mathbb{Z}_m^*$. It follows⁷ that for $k \in \mathbb{Z}_m^*$, $\tau_k(\mathfrak{q}_i) = \mathfrak{q}_{ik^{-1} \pmod m}$
38 and $\tau_k^{-1} = \tau_{k^{-1} \pmod m}$. Also, for $k \in \mathbb{Z}_m^*$, R and R^\vee are fixed by τ_k and so $\tau_k(R_q) = R_q$. Hence, if a is distributed
39 uniformly in R_q , then $\tau_k(a)$ is also distributed uniformly in R_q .

40 For $i \in \mathbb{Z}_m^*$, it can be shown that the quotient group $R^\vee/(\mathfrak{q}_i R^\vee)$ has cardinality q and the representatives
41 of the q distinct cosets can be taken to be the elements of the set $\{0, \dots, q-1\}$. The cardinality of the set
42 R_q^\vee is q^n . Using the Chinese Remainder Theorem (CRT), it can be shown that there is an isomorphism \mathfrak{J}
43 from R_q^\vee to $\bigoplus_{i \in \mathbb{Z}_m^*} (R^\vee/(\mathfrak{q}_i R^\vee))$. Further, \mathfrak{J} can be efficiently computed in both the forward and the backward

⁷The following fact is used to obtain $\tau_k(\mathfrak{q}_i) = \mathfrak{q}_{ik^{-1} \pmod m}$: for j such that $i \equiv jk \pmod m$, $\tau_k(\zeta - \omega^i) = \tau_k(\zeta - \omega^{jk}) = \zeta^k - \omega^{jk} = \zeta^k - \omega^j = (\zeta - \omega^j)x$, where $x = \zeta^{k-1} + \omega^j \zeta^{k-2} + \dots + \omega^{j(k-1)}$ is in R .

directions. For $i \in \mathbb{Z}_m^*$, let $w_i \in \{0, \dots, q-1\}$ represent a coset of $R^\vee / (\mathfrak{q}_i R^\vee)$. Given $(w_i)_{i \in \mathbb{Z}_m^*}$, it is possible to efficiently construct $w \in R_q^\vee$ such that the i -th component of $\mathfrak{I}(w)$ is represented by w_i . For the sake of notational convenience, we let w denote $\mathfrak{I}^{-1}((w_i)_{i \in \mathbb{Z}_m^*})$.

4.1 Intermediate problems

The “search to decision” reduction is obtained by composing several individual reductions between intermediate problems. The search ring-LWE problem requires finding s in R_q^\vee . The first step of the reduction is to show that it is sufficient to find any one of the components in the image of s under the isomorphism \mathfrak{I} . The relevant intermediate problem is the following.

Ring-LWE over \mathfrak{q}_i . For $i \in \mathbb{Z}_m^*$, the \mathfrak{q}_i -LWE $_{q, \leq \alpha}$ problem is the following. For $s \in R_q^\vee$ and a positive real number $r \leq \alpha$, given access to samples from $A_{s,r}$, the requirement is to find the i -th component of $\mathfrak{I}(s)$.

Let the representatives of \mathbb{Z}_m^* be chosen from the set $\{1, \dots, m-1\}$ with the usual ordering. For $i \in \mathbb{Z}_m^*$, let $i-$ denote the largest element in \mathbb{Z}_m^* which is less than i with the convention that $1-$ is taken to be 0.

The distribution $A_{s,r}^i$. For $i \in \mathbb{Z}_m^* \cup \{0\}$, $s \in R_q^\vee$ and a positive real number r , the distribution $A_{s,r}^i$ over $R_q \times \mathbb{T}$ is defined in the following manner. A sample from $A_{s,r}$ consists of a pair (a, \mathbf{b}) , where $a \in R_q$ and $\mathbf{b} = \sigma((a \cdot s)/q) + \mathbf{e} \bmod \sigma(R^\vee)$. A sample from $A_{s,r}^i$ is a sample from $A_{s,r}$ whose k -th component for $k \leq i$ has been randomised by adding a uniform random $h_k \in \{0, 1, \dots, q-1\}$ to the k -th component of $a \cdot s$, thereby hiding the information⁸.

A sample from $A_{s,r}^i$ hides information about s with respect to the factors \mathfrak{q}_k of $\langle q \rangle$ for $k \in \mathbb{Z}_m^*$ and $k \leq i$. For $i \in \mathbb{Z}_m^*$ or $i = 0$, as i increases from 0 to $m-1$, information about s is hidden in one more \mathfrak{q}_i -component than in the previous step. At the beginning, i.e. $i = 0$, all the components in the output of \mathfrak{I} carry information about s , while at the end, i.e., $i = m-1$, the element $a \cdot s + h$ is a uniform random element of R_q^\vee which is independent of both s and a . So for a sample (a, \mathbf{b}) drawn from $A_{s,r}^{m-1}$, a is uniform over R_q and \mathbf{b} is independent of a ; further, \mathbf{b} is the sum modulo $\sigma(R^\vee)$ of a uniform random element of $\sigma(R_q^\vee)/q$ and an element drawn from the distribution D_r . Consequently, a sample drawn from $A_{s,r}^{m-1}$ is almost uniform over $R_q \times \mathbb{T}$ (see Lemma 5.13 of [19]).

Variable width ring-DLWE relative to \mathfrak{q}_i . For $i \in \mathbb{Z}_m^*$ and a positive real number α , the ring-VWDLWE $_{q, \leq \alpha}^i$ problem is the following. Given access to $A_{s,r}^j$ for $s \in R_q^\vee$, positive real number $r \leq \alpha$ and $j \in \{i, i-\}$, the requirement is to find j . In other words, the solver must determine whether or not the i -th component of the distribution has been randomised.

Remark 6. *The letters VWD before LWE denote ‘variable width decision’. In Definition 5.8 of [19] this problem was denoted WDLWE, meaning worst-case decision LWE. In our view, the use of “worst-case” is inappropriate, whereas “variable width” indicates an important feature of the problem.*

The next step is to consider a fixed width version of the DLWE problem with respect to the ideal \mathfrak{q}_i .

Ring DLWE (fixed width) relative to \mathfrak{q}_i . Let $i \in \mathbb{Z}_m^*$ and $r_0 > 0$ be a real number. The ring-DLWE $_{q, r_0}^i$ problem is the following. Choose s uniformly at random from R_q^\vee . The requirement is to distinguish between inputs from A_{s, r_0}^{i-} and A_{s, r_0}^i . Formally, let \mathcal{D}_1 be an algorithm which takes as input a list \mathcal{T} of samples from $A_{s, r}^j$

⁸ For $i \in \mathbb{Z}_m^*$, let $\chi(i)$ be the following distribution over R_q^\vee . For $k \in \mathbb{Z}_m^*$, choose $h_k \in \{0, \dots, q-1\}$ as follows: $h_k = 0$ for $k > i$; and for $k \leq i$, the h_k ’s are chosen independently and uniformly; return $h = \mathfrak{I}^{-1}((h_k)_{k \in \mathbb{Z}_m^*})$.

A sample from $A_{s, r}^i$ is obtained as follows. For $i \in \mathbb{Z}_m^*$, let h be sampled from $\chi(i)$. Choose $(a, \mathbf{b}) \leftarrow A_{s, r}$ and output $(a, \mathbf{b} + \sigma(h)/q \bmod \sigma(R^\vee))$ as a sample from $A_{s, r}^i$; for $i = 0$, the distribution $A_{s, r}^0$ is defined to be $A_{s, r}$.

\mathcal{A}_5 solves ring-LWE $_{q,\leq\alpha}$ <ul style="list-style-type: none"> • has access to $A_{s,r}$-oracle • $r \leq \alpha$ and s are unknown to \mathcal{A}_5 • output s 	\mathcal{A}_6 solves \mathfrak{q}_i -LWE $_{q,\leq\alpha}$, $i \in \mathbb{Z}_m^*$ <ul style="list-style-type: none"> • has access to $A_{s,r}$-oracle • $r \leq \alpha$ and s are unknown to \mathcal{A}_5 • output the i-th component of $\mathfrak{I}(s)$
\mathcal{A}_7 solves ring-VWDLWE $_{q,\leq\alpha}^i$, $i \in \mathbb{Z}_m^*$ <ul style="list-style-type: none"> • has access to $A_{s,r}^j$, $s \in R_q^\vee$, $r \leq \alpha$, $j \in \{i, i-\}$ • output j 	
\mathcal{D}_1 solves ring-DLWE $_{q,r}^i$ <ul style="list-style-type: none"> • distinguishes between $A_{s,r}^i$ and $A_{s,r}^{i-}$ • $s \in R_q^\vee$, $r \leq \alpha$ 	\mathcal{D}_2 solves ring-DLWE $_{q,r}$ <ul style="list-style-type: none"> • distinguishes between $A_{s,r}$ and $U(R_q \times \mathbb{T})$ • $s \in R_q^\vee$, $r \leq \alpha$

Figure 2: Inputs and outputs of algorithms \mathcal{A}_5 to \mathcal{A}_7 and distinguishers \mathcal{D}_1 and \mathcal{D}_2 .

1 with $j \in \{i-, i\}$ and outputs a bit. For a fixed $s \in R_q^\vee$, let $p_{s,0}$ (resp. $p_{s,1}$) be the probability that \mathcal{D}_1 outputs 1
2 when \mathcal{T} consists of samples from A_{s,r_0}^{i-} (resp. A_{s,r_0}^i), where the probability is taken over all components of the
3 input other than s as well as the internal coin tosses of \mathcal{D}_1 . The advantage of the distinguisher is $|p_{s,0} - p_{s,1}|$. For
4 $\epsilon_1, \epsilon_2 \in (0, 1]$, we say that \mathcal{D}_1 is an (ϵ_1, ϵ_2) -distinguisher if \mathcal{D}_1 has advantage at least ϵ_2 for at least a proportion
5 ϵ_1 of the set of possible $s \in R_q^\vee$,

6 **Remark 7.** *The above definition is based on Definition 5.10 of [19]. The formulation is different from that of*
7 *Definition 5.10 and is in the form that is actually used in Lemma 5.16 and Theorem 3.6 of [19].*

8 The search ring-LWE to ring-DLWE reduction involves the following algorithms.

- \mathcal{A}_6 : an algorithm to solve \mathfrak{q}_i -LWE $_{q,\leq\alpha}$
- \mathcal{A}_7 : an algorithm to solve ring-VWDLWE $_{q,\leq\alpha}^i$
- \mathcal{D}_1 : a distinguisher for ring-DLWE $_{q,r}^i$
- \mathcal{D}_2 : a distinguisher for ring-DLWE $_{q,r}$

9 While going through the description below, it will be helpful to keep in mind the inputs and outputs of the
10 various algorithms which are shown in Figure 2.

11 \mathcal{A}_6 has access to an oracle which returns samples from $A_{s,r}$ for some unknown s and unknown $r \leq \alpha$; \mathcal{A}_7 has
12 access to an oracle $A_{s,r}^j$ for some unknown s , unknown $r \leq \alpha$ and j equal to either i or $i-$.

13 The overall reduction proceeds in several steps to construct an algorithm \mathcal{A}_5 to solve ring-LWE by using a
14 distinguisher \mathcal{D}_2 for ring-DLWE $_{q,r_0}$ as an oracle. In the context of cryptography, ring-DLWE $_{q,r_0}$ is the problem
15 whose solution breaks the cryptosystem, and \mathcal{A}_5 , in turn, is used by \mathcal{A}_0 via \mathcal{A}_1 , \mathcal{A}_2 , \mathcal{A}_3 , and \mathcal{A}_4 to solve SIVP $_\gamma$.

16 In the first step, \mathcal{A}_5 is constructed by using an algorithm \mathcal{A}_6 as an oracle to solve \mathfrak{q}_i -LWE $_{q,\leq\alpha}$ for some i .
17 The idea is to individually compute the components of $\mathfrak{I}(s)$ using \mathcal{A}_6 . Since \mathcal{A}_6 works for a particular i , the
18 automorphisms of the number field are used to ensure that the j -th component of $\mathfrak{I}(s)$ is transferred to the
19 i -th component so that \mathcal{A}_6 can be applied. In the next step, \mathcal{A}_7 is constructed by using an algorithm \mathcal{A}_7 as an
20 oracle to solve ring-VWDLWE $_{q,\leq\alpha}^i$. There are q possible x values of the i -th component. For each x , \mathcal{A}_6 modifies
21 the LWE samples in a manner such that if x is the correct value of $s \bmod \mathfrak{q}_i R^\vee$, then the samples are from the
22 distribution $A_{s,r}^{i-}$, while if x is not equal to $s \bmod \mathfrak{q}_i R^\vee$, then the samples are from the distribution $A_{s,r}^i$. The
23 oracle \mathcal{A}_7 can be used to determine which of these two cases occurs.

24 The biggest step in the reduction is the construction of \mathcal{A}_7 using a distinguisher \mathcal{D}_1 for ring-DLWE $_{q,r_0}^i$ as an
25 oracle. \mathcal{D}_1 is an oracle that, given two ℓ -tuples of samples that are randomised in the first i (resp. $i-$) components
26 of H and come from LWE-sampling with known Gaussian error distribution in the remaining components, can
27 distinguish between them. The construction of \mathcal{A}_7 shows that using \mathcal{D}_1 , one can answer the same question when

1 the LWE-sampling is with a Gaussian error distribution of width that is unknown (but less than a known bound).
2 The last step is the construction of \mathcal{D}_1 by using a distinguisher \mathcal{D}_2 for ring-DLWE $_{q,r_0}$ as an oracle.
3 Further details of the various steps are provided below with a focus on concrete aspects.

4 4.2 Reducing ring-LWE $_{q,\leq\alpha}$ to \mathfrak{q}_i -LWE $_{q,\leq\alpha}$

5 Suppose \mathcal{A}_6 is an algorithm to solve \mathfrak{q}_i -LWE $_{q,\leq\alpha}$ for some particular $i \in \mathbb{Z}_m^*$. We provide a brief description of
6 the construction of Algorithm \mathcal{A}_5 using \mathcal{A}_6 as an oracle (see Lemma 5.5 of [19] for details). The goal of \mathcal{A}_5
7 is to compute s . This can be done if each component of $\mathfrak{I}(s)$ can be computed. Algorithm \mathcal{A}_6 can compute
8 the i -th component of $\mathfrak{I}(s)$. \mathcal{A}_5 uses automorphisms to map the j -th component to the i -th component in the
9 following manner. For each $j \in \mathbb{Z}_m^*$, let $k = j \cdot i^{-1} \bmod m$. \mathcal{A}_5 then invokes \mathcal{A}_6 , and whenever \mathcal{A}_6 makes a
10 query for an LWE sample, \mathcal{A}_5 queries its own LWE oracle to obtain a sample (a, \mathbf{b}) . It responds to the query by
11 sending $(\tau_k(a), \tau_k(\mathbf{b}))$ to \mathcal{A}_6 . Let the output of \mathcal{A}_6 on the j -th invocation be denoted as s_j . Finally, \mathcal{A}_5 returns
12 $\mathfrak{I}^{-1}((s_j)_{j \in \mathbb{Z}_m^*})$.

13 **Proposition 4.** \mathcal{A}_5 invokes \mathcal{A}_6 a total of n times. The numbers of LWE queries made by \mathcal{A}_5 and \mathcal{A}_6 are equal.

14 It is interesting to note that the reduction of ring-LWE $_{q,\leq\alpha}$ to \mathfrak{q}_i -LWE $_{q,\leq\alpha}$ is made possible because ideal
15 lattices for cyclotomic number fields possess some nice algebraic properties. On the other hand, the reduction
16 itself can be considered to be a step in a possible attack on the search ring-LWE problem. This is because thanks
17 to the automorphisms, if an algorithm evaluates the i -th component for any fixed i , then it is possible to use it
18 to evaluate all the components; and for a fixed i the search space is not too big for exhaustive search.

19 The construction of \mathcal{A}_5 from \mathcal{A}_6 is based on two points, the existence of n automorphisms and the split of $\langle q \rangle$
20 into linear factors. There are two directions in which one could generalise from cyclotomic fields with primes q
21 that split completely. First, in a cyclotomic field one can take any prime q that does not divide m , in which case
22 $\langle q \rangle$ splits into n/f distinct prime ideals where f is the residue field degree. In other words, $R/(\mathfrak{q}_i R)$ is no longer
23 the field of q elements, but rather is a degree- f extension. In that case there would still be a contribution of n
24 to the tightness gap, because while there are only n/f prime ideals, for each prime ideal it would be required
25 to find f coordinates, so the Chinese Remainder Theorem presumably takes as much work as before (or more).
26 This approach has been briefly mentioned in footnote 8 on Page 26 of [19]. Second, one can generalize to non-
27 cyclotomic Galois fields. There it is still possible to find primes q that split completely, and if q does not split
28 completely, there will again be n/f prime ideals with residue field of degree f (provided q does not divide the
29 discriminant of the field).

30 4.3 Reducing \mathfrak{q}_i -LWE $_{q,\leq\alpha}$ to ring-VWDLWE $_{q,\leq\alpha}^i$

31 Suppose \mathcal{A}_7 is an algorithm to solve ring-VWDLWE $_{q,\leq\alpha}^i$. This algorithm is used as an oracle to construct an
32 algorithm \mathcal{A}_6 to solve \mathfrak{q}_i -LWE $_{q,\leq\alpha}$. We provide a brief description of the construction and for details we refer to
33 Lemma 5.9 of [19]. The requirement for \mathcal{A}_6 is to determine the i -th component of $\mathfrak{I}(s)$. As mentioned earlier, each
34 component of $\mathfrak{I}(s)$ can be represented by an element from the set $\{0, \dots, q-1\}$. For each $x \in \{0, \dots, q-1\}$, \mathcal{A}_6
35 does the following. It first computes an element $g \in R_q^\vee$ such that $\mathfrak{I}(g)$ is equal to x in the i -th component and is
36 equal to zero in all other components. Then \mathcal{A}_6 invokes \mathcal{A}_7 . For each LWE query made by \mathcal{A}_7 , \mathcal{A}_6 queries its own
37 oracle to obtain a pair (a, \mathbf{b}) . It then computes an element $v \in R_q$ such that under the isomorphism from R_q to
38 $\bigoplus_{j \in \mathbb{Z}_m^*} R/(\mathfrak{q}_j R)$, the i -th component is chosen uniformly at random from $\{0, \dots, q-1\}$ and all other components
39 are equal to zero. Next, \mathcal{A}_6 adds $(v, \sigma((v \cdot g)/q) \bmod \sigma(R^\vee))$ to the sample (a, \mathbf{b}) and then randomises the first
40 i - components by adding a random element to the second part of each of those components of the sample. \mathcal{A}_6
41 sends the resulting modified sample to \mathcal{A}_7 as its response to the query⁹. At the end of its oracle queries, if \mathcal{A}_7
42 returns i -, then \mathcal{A}_6 returns x . The crucial point for correctness is that if the value of x is equal to the i -th

⁹In other words, \mathcal{A}_6 generates h from the distribution $\chi(i-)$ and returns the pair $(a + v, \mathbf{b} + (\sigma(h + v \cdot g)/q) \bmod \sigma(R^\vee))$ to \mathcal{A}_7 .

1 component of $\mathcal{J}(s)$, then the samples returned to \mathcal{A}_7 are from the distribution $A_{s,r}^{i-}$, and if not, then the samples
 2 returned to \mathcal{A}_7 are from the distribution $A_{s,r}^i$.

3 **Proposition 5.** \mathcal{A}_6 invokes \mathcal{A}_7 at most q times. The numbers of LWE queries made by \mathcal{A}_6 and \mathcal{A}_7 are equal.

4 4.4 Reducing ring-VWDLWE $_{q,\leq\alpha}^i$ to ring-DLWE $_{q,\tau}^i$

5 For a fixed $i \in \mathbb{Z}_m^*$, Lemma 5.16 of [19] states¹⁰ that ring-VWDLWE $_{q,\leq\alpha}^i$ reduces to ring-DLWE $_{q,\tau}^i$ in randomised
 6 polynomial time, where

$$\tau = \alpha \cdot \left(\frac{n\ell}{\ln(n\ell)} \right)^{1/4} \quad (10)$$

7 and ℓ is the number of LWE samples required by the distinguisher for ring-DLWE $_{q,\tau}^i$. Note that (10) shows that
 8 $\tau > \alpha$. Since our goal is to estimate feasibility, we perform a concrete analysis which turns out to be considerably
 9 more complicated than the sketch of a proof in the asymptotic setting that was provided in [19]. In particular, in
 10 order to get a rigorous proof for this part of the reduction we needed to include the factor N_2 , defined below. This
 11 has a substantial effect on the ratio τ/α , which now becomes

$$\frac{\tau}{\alpha} = \left(\frac{nN_2\ell}{\ln(nN_2\ell)} \right)^{1/4}, \quad (11)$$

12 and that, in turn, brings a new term $N_2^{1/4} > \sqrt{n}$ (see below) into the SIVP approximation factor.

13 The essential difference between the problems ring-VWDLWE $_{q,\leq\alpha}^i$ and ring-DLWE $_{q,\tau}^i$ is in the distribution of
 14 the error of the LWE samples. For the former problem the errors follow D_r , where $r \leq \alpha$, while in the latter
 15 problem the errors follow D_τ with $\tau > \alpha$ as mentioned above. The reduction of ring-VWDLWE $_{q,\leq\alpha}^i$ to ring-DLWE $_{q,\tau}^i$
 16 is a trade-off between narrower width with no knowledge of the width except for an upper bound versus wider
 17 width (which generally means less useful samples) with knowledge of the width.

18 Let $\mathcal{T} = ((a_k, \mathbf{b}_k))_{1 \leq k \leq \ell}$ be a list of ℓ samples from $A_{s,r}^j$, where $r \leq \alpha$ and j is equal to either $i-$ or i . Suppose
 19 $t \in R_q^\vee$ and let $\mathbf{f}_1, \dots, \mathbf{f}_\ell$ be chosen independently from D_τ . Define a list $\mathcal{T}' = ((a'_k, \mathbf{b}'_k))_{1 \leq k \leq \ell}$, where $a'_k = a_k$ and
 20 $\mathbf{b}'_k = \mathbf{b}_k + \sigma(a_k \cdot t)/q + \mathbf{f}_k \pmod{\sigma(R^\vee)}$. Since (a_k, \mathbf{b}_k) is a sample from $A_{s,r}^j$, the error vector \mathbf{e}_k in \mathbf{b}_k follows
 21 D_r . The error in \mathbf{b}'_k , which is $\mathbf{e}_k + \mathbf{f}_k$, follows $D_{r'}$, where $r' = \sqrt{r^2 + \tau^2}$. Hence, the samples in \mathcal{T}' are from the
 22 distribution $A_{s+t,r'}^j$.

23 From \mathcal{T}' , a list \mathcal{T}'' is obtained as follows. For each pair (a'_k, \mathbf{b}'_k) in \mathcal{T}' , a'_k is unchanged and \mathbf{b}'_k is modified
 24 so as to partially randomise \mathbf{b}_k ; namely all components of $a_k \cdot (s+t)$ up through the i -th are randomised¹¹.
 25 Note that irrespective of the value of j , the samples in \mathcal{T}'' are from the distribution $A_{s+t,r'}^i$. Note that the LWE
 26 secret in both \mathcal{T}' and \mathcal{T}'' is $s+t$ and the distribution of the LWE errors is $D_{r'}^\ell$. This ensures two things. The
 27 first is a random self-reduction where the LWE secret s is mapped to $s+t$ and the second is the addition of the
 28 \mathbf{f} -errors so as to get the error width r' to within a small multiplicative factor of τ .

29 Let \mathcal{D}_1 be an (ϵ_1, ϵ_2) -distinguisher for ring-DLWE $_{q,\tau}^i$. Using \mathcal{D}_1 , an algorithm \mathcal{A}_7 for ring-VWDLWE $_{q,\leq\alpha}^i$ is
 30 constructed as follows. \mathcal{A}_7 has access to an oracle for $A_{s,r}^j$, where $r \leq \alpha$ and j is either i or $i-$. The requirement
 31 is to determine j . The construction of \mathcal{A}_7 has two nested loops: an outer loop of N_1 iterations and for each of
 32 these iterations an inner loop of N_2 iterations. In each iteration of the outer loop, \mathcal{A}_7 chooses t uniformly at
 33 random from R_q^\vee . Then the inner loop of N_2 iterations starts. In each of the N_2 iterations of the inner loop,

¹⁰The actual statement of Lemma 5.16 in [19] is for the case where the error distribution for the ring-VWDLWE problem is from a family of elliptical Gaussian distributions. Here we consider the simpler situation where the error distribution is from a family of spherical Gaussian distributions. This simplification does not have any effect on the concrete security analysis.

¹¹In other words, (a'_k, \mathbf{b}'_k) is modified to $(a'_k, \mathbf{b}'_k + \sigma(h_k)/q \pmod{\sigma(R^\vee)})$, where h_k is chosen from $\chi(i)$.

1 \mathcal{A}_7 obtains a list \mathcal{T} of samples from $A_{s,r}^j$; chooses $\mathbf{f}_1, \dots, \mathbf{f}_\ell$ independently from $D_{\mathfrak{t}}^\ell$; and uses \mathcal{T} , t and $\mathbf{f}_1, \dots, \mathbf{f}_\ell$
2 to prepare the lists \mathcal{T}' and \mathcal{T}'' as described above. Then \mathcal{A}_7 runs \mathcal{D}_1 on \mathcal{T}' and \mathcal{T}'' obtaining in return the
3 corresponding 1-bit outputs. At the end of the inner loop of N_2 iterations, \mathcal{A}_7 obtains estimates $\hat{\mathbf{p}}_0$ and $\hat{\mathbf{p}}_1$ of
4 the probabilities \mathbf{p}_0 and \mathbf{p}_1 that \mathcal{D}_1 accepts inputs from the distributions $A_{s+t,r'}^j$ and $A_{s+t,r'}^i$ respectively. If
5 $|\hat{\mathbf{p}}_0 - \hat{\mathbf{p}}_1| \geq \epsilon_2/4$, then \mathcal{A}_7 returns $i-$ and stops. If in none of the N_1 iterations the condition $|\hat{\mathbf{p}}_0 - \hat{\mathbf{p}}_1| \geq \epsilon_2/4$ is
6 satisfied, then \mathcal{A}_7 returns i . Figure 2 in Appendix B provides a pseudo-code description of \mathcal{A}_7 .

7 Note that \mathcal{D}_1 is supposed to work for errors following the distribution $D_{\mathfrak{t}}^\ell$. It is, however, invoked on the lists
8 \mathcal{T}' and \mathcal{T}'' , where the errors in the samples in these lists follow $D_{r'}^\ell$. Due to the change in the width of the error
9 distribution, \mathcal{D}_1 may not behave as an (ϵ_1, ϵ_2) -distinguisher. This is taken care of in the following analysis.

10 Algorithm \mathcal{A}_7 fails if it returns an incorrect answer. This can happen in two ways, namely that $j = i$ and \mathcal{A}_7
11 returns $i-$, and $j = i-$ and \mathcal{A}_7 returns i . We call the former to be Type-1 failure and the latter to be Type-2
12 failure.

13 **Case $j = i$.** In this case, the samples in both \mathcal{T}' and \mathcal{T}'' follow $A_{s+t,r'}^i$ and so $\mathbf{p}_0 = \mathbf{p}_1$. Consider any one of the
14 N_1 iterations of the outer loop. From the additive form of the Chernoff-Hoeffding bound [16], we have

$$\begin{aligned} \Pr[\mathbf{p}_0 - \epsilon_2/8 \leq \hat{\mathbf{p}}_0 \leq \mathbf{p}_0 + \epsilon_2/8] &\geq 1 - 2 \exp(-N_2 \epsilon_2^2/32), \\ \Pr[\mathbf{p}_1 - \epsilon_2/8 \leq \hat{\mathbf{p}}_1 \leq \mathbf{p}_1 + \epsilon_2/8] &\geq 1 - 2 \exp(-N_2 \epsilon_2^2/32). \end{aligned}$$

16 Since $\mathbf{p}_0 = \mathbf{p}_1$, it follows that $\Pr[|\hat{\mathbf{p}}_0 - \hat{\mathbf{p}}_1| \leq \epsilon_2/4] \geq 1 - 4 \exp(-N_2 \epsilon_2^2/32)$. So in any of the N_1 iterations of the
17 outer loop the probability of Type-1 failure is at most $4 \exp(-N_2 \epsilon_2^2/32)$. The probability of Type-1 failure in any
18 of the N_1 iterations is then at most $4N_1 \exp(-N_2 \epsilon_2^2/32)$.

19 **Case $j = i-$.** In this case, the samples in \mathcal{T}' follow $A_{s+t,r'}^{i-}$ while the samples in \mathcal{T}'' follow $A_{s+t,r'}^i$. In any of
20 the N_2 iterations of the inner loop, let $\mathbf{z}_1 = \mathbf{e}_1 + \mathbf{f}_1, \dots, \mathbf{z}_\ell = \mathbf{e}_\ell + \mathbf{f}_\ell$ be the errors in the LWE samples in the
21 lists \mathcal{T}' and \mathcal{T}'' . Let \mathbf{z} be a vector consisting of all the $N_2 \ell$ errors in the N_2 iterations of the loop. The vector \mathbf{z}
22 follows $D_{r'}^{\ell N_2}$. Suppose instead that it follows $D_{\mathfrak{t}}^{\ell N_2}$. Later we will compute a correction factor to account for the
23 width being r' rather than \mathfrak{t} . We denote the corresponding probabilities and their estimates by p_0, p_1, \hat{p}_0 and
24 \hat{p}_1 . Let $p_{s+t,0}$ and $p_{s+t,1}$ respectively denote the probabilities p_0 and p_1 corresponding to a particular value of t .
25 Similarly, let $\hat{p}_{s+t,0}$ and $\hat{p}_{s+t,1}$ respectively denote the estimates \hat{p}_0 and \hat{p}_1 corresponding to a particular value of
26 t . Further, let $\hat{p}_{s+t,\mathbf{z},0}$ and $\hat{p}_{s+t,\mathbf{z},1}$ denote these estimates for a particular value of t and \mathbf{z} .

28 We say that a value $s+t$ is good if $|p_{s+t,0} - p_{s+t,1}| \geq \epsilon_2$. From the definition of an (ϵ_1, ϵ_2) -distinguisher, the
29 probability of a good $s+t$ is at least ϵ_1 . For a good $s+t$, using the additive form of the Chernoff-Hoeffding
30 bound [16], we have

$$\begin{aligned} \Pr[p_{s+t,0} - \epsilon_2/4 \leq \hat{p}_{s+t,0} \leq p_{s+t,0} + \epsilon_2/4] &\geq 1 - 2 \exp(-N_2 \epsilon_2^2/8), \\ \Pr[p_{s+t,1} - \epsilon_2/4 \leq \hat{p}_{s+t,1} \leq p_{s+t,1} + \epsilon_2/4] &\geq 1 - 2 \exp(-N_2 \epsilon_2^2/8). \end{aligned}$$

31 Since the events $p_{s+t,0} - \epsilon_2/4 \leq \hat{p}_{s+t,0} \leq p_{s+t,0} + \epsilon_2/4$ and $p_{s+t,1} - \epsilon_2/4 \leq \hat{p}_{s+t,1} \leq p_{s+t,1} + \epsilon_2/4$ along with the
32 condition $|p_{s+t,0} - p_{s+t,1}| \geq \epsilon_2$ together imply $|\hat{p}_{s+t,0} - \hat{p}_{s+t,1}| \geq \epsilon_2/2$, we obtain

$$\Pr[|\hat{p}_{s+t,0} - \hat{p}_{s+t,1}| \geq \epsilon_2/2] \geq 1 - 4 \exp(-N_2 \epsilon_2^2/8). \quad (12)$$

33 For N_2 about ϵ_2^{-2} times a constant, the difference $|\hat{p}_{s+t,0} - \hat{p}_{s+t,1}|$ will be at least $\epsilon_2/2$ with probability almost 1.
34 Keeping this in mind, henceforth we will assume (see Remark 8 below)

$$N_2 = \tilde{O}(\epsilon_2^{-2}). \quad (13)$$

35 For simplicity, we also assume that for N_2 given by (13), the following holds.

$$\text{If } s+t \text{ is good, then } |\hat{p}_{s+t,0} - \hat{p}_{s+t,1}| \geq \epsilon_2/2 \text{ holds with probability 1.} \quad (14)$$

1 Given a good $s + t$, we say that \mathbf{z} is good if $|\hat{p}_{s+t,\mathbf{z},0} - \hat{p}_{s+t,\mathbf{z},1}| \geq \epsilon_2/4$ holds. Using Proposition 10 in
2 Appendix B, the probability of a good \mathbf{z} (for a good $s + t$) is at least $\epsilon_2/4$. Changing the error distribution
3 from $D_{\mathbf{r}}^{\ell N_2}$ to $D_{\mathbf{r}'}^{\ell N_2}$, the probability of a good \mathbf{z} under $D_{\mathbf{r}'}^{\ell N_2}$ (i.e., one for which $|\hat{p}_{s+t,\mathbf{z},0} - \hat{p}_{s+t,\mathbf{z},1}| \geq \epsilon_2/4$ holds)
4 is at least $\epsilon_2^2/(256nN_2\ell)^{1/2}$ (see Proposition 14 in Appendix B). So the probability of a good pair $(s + t, \mathbf{z})$
5 where \mathbf{z} follows $D_{\mathbf{r}'}^{\ell N_2}$ is at least $\epsilon_1\epsilon_2^2/(256nN_2\ell)^{1/2}$. Consequently, if N_1 is around $(256nN_2\ell)^{1/2}/(\epsilon_1\epsilon_2^2)$, then with
6 probability exponentially close to 1 a good tuple will be encountered in one of the iterations of the outer loop.

7 Type-2 failure can occur in two ways. The first way is that in none of the N_1 iterations, a good tuple is
8 obtained. The second way is that for a good tuple, the condition $|\hat{p}_{s+t,\mathbf{z},0} - \hat{p}_{s+t,\mathbf{z},1}| \geq \epsilon_2/4$ does not hold. The
9 above analysis shows that the probability of either of these errors is exponentially small.

10 The number of times \mathcal{A}_7 calls \mathcal{D}_1 is N_1N_2 which is about $(256n\ell)^{1/2}N_2^{3/2}/(\epsilon_1\epsilon_2^2)$ and the number of samples
11 of $A_{s,r}^j$ required by \mathcal{A}_7 is $N_1N_2\ell$. As mentioned above N_2 is about ϵ_2^{-2} .

12 **Remark 8.** Ignoring Type-2 failures, the probability that \mathcal{A}_7 fails is given by the probability of Type-1 fail-
13 ure. As shown above, this probability is at most $4N_1 \exp(-N_2\epsilon_2^2/32)$. In the complete reduction of K-SIVP $_{\gamma}$ to
14 ring-VWDLWE $_{q,\leq\alpha}^i$, let N_3 be the number of times \mathcal{A}_7 is called. Then the probability that any of these calls fails is
15 at most $4N_1N_3 \exp(-N_2\epsilon_2^2/32)$. This shows that the value of N_2 should be about $32L\epsilon_2^{-2}$ where, $L = \ln N_1 + \ln N_3$.
16 Since N_2^2 appears in the number of oracle calls, the factor $32L$ contributes $2^{10}L^2$ to the tightness gap. For practical
17 values of the parameters, this can be significant.

18 In view of the above, we have the following result.

19 **Proposition 6.** The number of times \mathcal{A}_7 calls \mathcal{D}_1 is about $(\epsilon_1\epsilon_2^5)^{-1} \cdot (256n\ell)^{1/2}$. The number of times \mathcal{A}_7 calls
20 its LWE oracle is about $(\epsilon_1\epsilon_2^5)^{-1} \cdot (256n)^{1/2}\ell^{3/2}$.

21 **Remark 9.** We tried various ways to fill in the details of this reduction while maintaining (10) rather than
22 resorting to the much larger \mathbf{r}/α ratio in (11). The factor (11) came from the need to bound the effect on the
23 probability measure of replacing the distribution $D_{\mathbf{r}}^{N_2\ell}$ by the actual distribution $D_{\mathbf{r}'}^{N_2\ell}$ that governs the error-
24 vectors \mathbf{z} that go into a set of N_2 pairs of lists $(\mathcal{T}', \mathcal{T}'')$ that are input to \mathcal{D}_2 .

25 One promising possibility seemed to be that we could avoid the N_2 term if we were able to consider each inner
26 loop to have an N_2 -tuple of lists coming from a fixed set of ℓ samples from $A_{s+t,r}^j$ with fixed t and fixed added
27 error-vectors \mathbf{f} , so that the only quantities that vary within an inner loop would be the randomisation h 's and the
28 internal coin tosses of the distinguisher.

29 However, in the case $j = i$ we would need to know that, for any fixed $s + t$ and for any fixed ℓ -tuple of \mathbf{a} and \mathbf{z}
30 coming from the ℓ samples from $A_{s+t,r}^j$, the probabilities of output 1 for \mathcal{T}' and for \mathcal{T}'' are equal. This is true if
31 the distributions of h 's are the same for the two lists. However, they aren't: for \mathcal{T}' the distribution is $\chi(i-)$ and
32 for \mathcal{T}'' it's $\chi(i)$. Speaking less formally, the difference is that the i -th component of \mathcal{T}' is randomised but fixed
33 for all of the N_2 lists, while the i -th component of \mathcal{T}'' gets a new randomisation in each new list. One can adopt
34 the heuristic assumption that the distinguisher sees them as indistinguishable, but in that case we don't have a
35 rigorous argument.

36 **Remark 10.** It is possible to use (10), but in that case we obtain a super-exponential time algorithm. See
37 Remark 23 in Appendix B.

38 4.5 Reducing ring-DLWE $_{q,\mathbf{r}}^i$ to ring-DLWE $_{q,\mathbf{r}}$

39 The final requirement is a reduction to the decision problem ring-DLWE (see the proof of Lemma 5.14 of [19]).
40 Suppose \mathcal{D}_2 is a (δ_1, δ_2) -distinguisher for ring-DLWE $_{q,\mathbf{r}}$, i.e., it has advantage at least δ_2 on a fraction δ_1 of the set
41 of possible values of s . We show that there is an $i \in \mathbb{Z}_m^*$, such that \mathcal{D}_2 will function as a $(\delta_1/n, \delta_2/n)$ -distinguisher
42 \mathcal{D}_1 for ring-DLWE $_{q,\mathbf{r}}^i$.

1 We say that s is “useful” if \mathcal{D}_2 has advantage at least δ_2 in distinguishing $A_{s,\mathfrak{r}}$ from the uniform distribution
2 over $R_q \times \mathbb{T}$. From the assumption on \mathcal{D}_2 , the proportion of useful $s \in R_q^\vee$ is at least δ_1 . For $s \in R_q^\vee$ and
3 $i \in \mathbb{Z}_m^* \cup \{0\}$, let $p_{s,i}$ be the probability that \mathcal{D}_2 outputs 1 on being provided ℓ samples from $A_{s,\mathfrak{r}}^i$. Recall
4 that $A_{s,\mathfrak{r}} = A_{s,\mathfrak{r}}^0$ and that samples from $A_{s,\mathfrak{r}}^{m-1}$ are almost uniformly distributed over $R_q \times \mathbb{T}$. So for a useful s ,
5 $|p_{s,0} - p_{s,m-1}| \geq \delta_2$. It then follows that there is at least one $i \in \mathbb{Z}_m^*$ such that $|p_{s,i-} - p_{s,i}| \geq \delta_2/n$. We say
6 that a pair $(s, i) \in R_q^\vee \times \mathbb{Z}_m^*$ is “useful” if the condition $|p_{s,i-} - p_{s,i}| \geq \delta_2/n$ holds. Since the proportion of useful
7 $s \in R_q^\vee$ is at least δ_1 , and for each useful s there is at least one i such that (s, i) is useful, it follows that for
8 some i there is a proportion at least δ_1/n of s values such that (s, i) is useful. We then fix such an i , denoted
9 i_0 . We define s to be “good” if (s, i_0) is useful. We attempt to use \mathcal{D}_2 as our \mathcal{D}_1 for the n possible values of
10 i , and when we get to $i = i_0$ we will find that our (δ_1, δ_2) -distinguisher \mathcal{D}_2 for ring-DLWE $_{q,\mathfrak{r}}$ also functions as a
11 $(\delta_1/n, \delta_2/n)$ -distinguisher for ring-DLWE $_{q,\mathfrak{r}}^i$.

12 **Proposition 7.** *Suppose \mathcal{D}_2 is a (δ_1, δ_2) -distinguisher for ring-DLWE $_{q,\mathfrak{r}}$. Then there is some $i \in \mathbb{Z}_m^*$ such that
13 there is an (ϵ_1, ϵ_2) -distinguisher \mathcal{D}_1 for ring-DLWE $_{q,\mathfrak{r}}^i$, where $\epsilon_1 = \delta_1/n$ and $\epsilon_2 = \delta_2/n$.*

14 Since $\epsilon_2 = \delta_2/n$, the value of N_2 given by (13) becomes

$$N_2 = \tilde{O}(n^2 \delta_2^{-2}). \quad (15)$$

15 **Remark 11.** *Proposition 7 does not show how to choose an i for which \mathcal{D}_1 exists. The argument showing
16 existence of i does not help determine the value of i . We can repeat the full reduction of K -SIVP $_\gamma$ to ring-
17 DLWE $_{q,\mathfrak{r}}$ for each possible i , and among the responses we select the basis that has the smallest maximum length.
18 Since there are n possibilities, this introduces a factor of n into the tightness gap.*

19 4.6 The tightness gap in the search ring-LWE to ring-DLWE reduction

20 The analysis in this section showed how to construct an algorithm \mathcal{A}_5 to solve ring-LWE $_{q,\leq\alpha}$ using a (δ_1, δ_2) -
21 distinguisher \mathcal{D}_2 for ring-DLWE $_{q,\mathfrak{r}}$. This required going through algorithms \mathcal{A}_6 and \mathcal{A}_7 and the distinguisher
22 \mathcal{D}_1 .

23 First we consider the number of times \mathcal{A}_5 calls \mathcal{D}_2 . By Proposition 4, \mathcal{A}_5 calls \mathcal{A}_6 a total of n times and
24 by Proposition 5, \mathcal{A}_6 calls \mathcal{A}_7 a total of q times. By Proposition 6, \mathcal{A}_7 calls \mathcal{D}_1 about $(\epsilon_1 \epsilon_2^5)^{-1} \cdot (256n\ell)^{1/2}$
25 times. From the discussion in Section 4.5, \mathcal{D}_1 is identical to \mathcal{D}_2 . Thus, the number of times \mathcal{A}_5 calls \mathcal{D}_2 is about
26 $qn(\epsilon_1 \epsilon_2^5)^{-1} \cdot (256n\ell)^{1/2}$.

27 Next we consider the number of LWE queries made by \mathcal{A}_5 . By Propositions 4 and 5, the number of LWE
28 queries made by \mathcal{A}_5 , \mathcal{A}_6 and \mathcal{A}_7 are equal. By Proposition 6, the number of LWE queries made by \mathcal{A}_7 is about
29 $(\epsilon_1 \epsilon_2^5)^{-1} \cdot ((256n)^{1/2} \ell^{3/2})$. Hence, the number of LWE queries made by \mathcal{A}_5 is also $(\epsilon_1 \epsilon_2^5)^{-1} \cdot ((256n)^{1/2} \ell^{3/2})$.

30 By Proposition 7, we have $\epsilon_1 = \delta_1/n$ and $\epsilon_2 = \delta_2/n$. Using these values of ϵ_1 and ϵ_2 in the above expressions,
31 we obtain the following result.

32 **Theorem 8.** *The number of times \mathcal{A}_5 calls \mathcal{D}_2 is about $qn(\delta_1 \delta_2^5)^{-1} n^6 \cdot (256n\ell)^{1/2} \approx qn^{15/2} \ell^{1/2} (\delta_1 \delta_2^5)^{-1}$ and the
33 number of LWE queries made by \mathcal{A}_5 is about $(\delta_1 \delta_2^5)^{-1} n^6 \cdot ((256n)^{1/2} \ell^{3/2}) \approx n^{13/2} \ell^{3/2} (\delta_1 \delta_2^5)^{-1}$.*

34 4.7 The parameters γ and q

35 In order to have confidence in the security of a ring-DLWE-based cryptosystem, we want to be sure that unless
36 (δ_1, δ_2) is negligible, there is no efficient (δ_1, δ_2) -distinguisher for ring-DLWE $_{q,\mathfrak{r}}$, where q (the modulus) and \mathfrak{r} (the
37 distribution width) are parameters of our cryptosystem. In particular, we don’t want there to be an efficient
38 distinguisher with advantage $2^{-\beta_2}$ unless β_2 is fairly large. We want the K -SIVP $_\gamma$ problem that reduces to ring-
39 DLWE $_{q,\mathfrak{r}}$ with this choice of δ_1, δ_2 to be hard. Below we investigate this question using the value $n = 2^{10}$, which
40 has been given as a parameter for some proposed cryptosystems [1, 6].

δ_2	2^{-128}	2^{-50}	2^{-25}
$\gamma = 2^k$	$2^{81.5}$	$2^{42.5}$	2^{30}
$2^{n/k}$	2^{12}	2^{24}	2^{34}
q	2^{76}	2^{37}	2^{34}

Table 1: For $n = 2^{10}$ the lower bounds on γ and upper bounds on $2^{n/k}$ along with lower bounds on q .

According to Theorem 3, the approximation factor γ in the SIVP is \tilde{O} of the following expression:

$$\frac{\sqrt{n}}{\alpha} = \frac{\sqrt{n}}{\mathfrak{r}} \left(\frac{nN_2\ell}{\ln(nN_2\ell)} \right)^{1/4} \quad (16)$$

by (11). Note that the parameter δ_1 does not have any effect on γ . As explained in the last paragraph of §3.1 of [19], in the case of a cyclotomic ideal lattice with n a power of 2 we need \mathfrak{r} to be bounded above by $O(\sqrt{\log n/n})$, or else the distribution will be statistically indistinguishable from uniform and no distinguisher will be possible.

Using (15) and the bound $O(\sqrt{\log n/n})$ on \mathfrak{r} and ignoring log factors and constants, from (16) we have

$$\gamma > n(nN_2\ell)^{1/4} > n(n\ell n^2\delta_2^{-2})^{1/4} = n^{7/4}\ell^{1/4}\delta_2^{-1/2} > n^{7/4}\delta_2^{-1/2}. \quad (17)$$

For example, choosing $n = 2^{10}$ and $\delta_2 = 2^{-\beta_2}$, we find that $\gamma > 2^{(35+\beta_2)/2}$.

Now for $\gamma = 2^k$ the fastest classical algorithm known that solves SVP_γ (and hence also solves $K\text{-SVP}_\gamma$ and its equivalent $K\text{-SIVP}_\gamma$) has running time $2^{\tilde{\theta}(n/k)}$ where $\tilde{\theta}$ suppresses a log factor [24]. We clearly want $2^{n/k}$ to be large. From $\gamma = 2^k > 2^{(35+\beta_2)/2}$, we have $k > (35 + \beta_2)/2$. Suppose we are considering 128-bit security. If we are extra cautious, then we will choose $\beta_2 = 128$; if we are less cautious, then we may choose $\beta_2 = 50$; and if we are not particularly risk-averse we might choose $\beta_2 = 25$. The corresponding lower bounds on γ and upper bounds on $2^{n/k}$ are shown in Table 1. None of these values inspire confidence in the hardness of $K\text{-SIVP}_\gamma$. In particular, the approximation factors γ are very large, and the running times $2^{n/k}$ are too small.

In addition, a practicality issue arises when we consider the modulus q . A condition for the reductions is that $q\alpha > 2\omega(\sqrt{\ln n})$. Using (11), we obtain

$$q > \frac{2\omega(\sqrt{\ln n})}{\alpha} = \frac{2\omega(\sqrt{\ln n})}{\mathfrak{r}} \left(\frac{nN_2\ell}{\ln(nN_2\ell)} \right)^{1/4}. \quad (18)$$

Again ignoring constants and log-terms and using $\mathfrak{r} < O(\sqrt{\log n/n})$, we have

$$q > n^{5/4}\ell^{1/4}\delta_2^{-1/2} > n^{5/4}\delta_2^{-1/2}. \quad (19)$$

With our values $n = 2^{10}$, $\delta_2 = 2^{-\beta_2}$ we find that $q > 2^{(25+\beta_2)/2}$. The lower bounds for q corresponding to $\beta_2 = 128, 50$ and 25 are shown in Table 1. The cryptosystem would be quite inefficient with these values of the modulus.

Remark 12. *The NIST-PQC proposals SABER and Kyber are based on module lattices. In §7.1 we consider the reduction for module lattices. The designers of Kyber chose parameters based on ideal lattices of the same dimension as that of module lattices. (See the first paragraph on the second page of [7].) We make the same assumptions for SABER. There are several variants of SABER and Kyber, and the highest security variant for both sets the dimension $n = 1024$. Irrespective of the dimension, the values¹² of q for SABER and Kyber are 2^{13} and 3329 respectively. These values are much lower than the values of q in Table 1.*

¹²Downloaded from <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions> on February 28, 2022.

5 The tightness gap in the K -SIVP $_{\gamma}$ to ring-DLWE $_{q,\tau}$ reduction

The tightness gap for the entire reduction is the number of times \mathcal{A}_0 calls \mathcal{D}_2 . This is given by the product of the number of times \mathcal{A}_0 calls \mathcal{A}_5 and the number of times \mathcal{A}_5 calls \mathcal{D}_2 . The former is given by Theorem 3 to be $(2n + (\log n)/2)n^4 \cdot N$, where N is the number of LWE queries made by \mathcal{A}_5 , and from Theorem 8 the latter is about $qn^{15/2}\ell^{1/2}(\delta_1\delta_2^5)^{-1}$. Also from Theorem 8, N is about $n^{13/2}\ell^{3/2}(\delta_1\delta_2^5)^{-1}$. Remark 11 shows that we have an additional factor of n in the number of times \mathcal{A}_0 calls \mathcal{D}_2 .

Based on the above analysis, the concrete version of the complete reduction of K -SIVP $_{\gamma}$ to ring-DLWE $_{q,\tau}$ is given by the following theorem, which corresponds to Theorem 3.6 of [19].

Theorem 9. *Let K be the m -th cyclotomic number field having degree $n = \varphi(m)$ and $R = \mathcal{O}_K$ be its ring of integers. Let τ be a positive real number bounded from above by $O(\sqrt{\log n/n})$. Let $\delta_1, \delta_2 \in (0, 1]$. Let q be a prime greater than 2 such that $q \equiv 1 \pmod{m}$ and $q > (2\omega(\sqrt{\ln n})/\tau) \cdot (nN_2\ell/\ln(nN_2\ell))^{1/4}$, where N_2 is defined in the course of the proof and has magnitude $\tilde{O}(n^2/\delta_2^2)$, and ℓ is a positive integer. Suppose there is a (δ_1, δ_2) -distinguisher \mathcal{D}_2 which solves ring-DLWE $_{q,\tau}$ given ℓ samples. Then there is a quantum algorithm \mathcal{A}_0 requiring approximately $3n^2$ logical qubits to solve K -SIVP $_{\gamma}$, where $\gamma = \tilde{O}(n^{5/4}\ell^{1/4}/(\tau\delta_2^{1/2}))$. The number of times \mathcal{A}_0 calls \mathcal{D}_2 is about*

$$(2n + (\log n)/2)n^4 \cdot n^{13/2}\ell^{3/2}(\delta_1\delta_2^5)^{-1} \cdot qn^{15/2}\ell^{1/2}(\delta_1\delta_2^5)^{-1} \cdot n \approx qn^{20}\ell^2 \cdot (\delta_1\delta_2^5)^{-2}. \quad (20)$$

From the point of view of practical cryptography, it is of interest to consider the tightness gap G for practical values of n . Let us consider $n = 2^{10}$ as in §4.7. Suppose $\delta_1 = 2^{-\beta_1}$, $\delta_2 = 2^{-\beta_2}$. Then q is at least $2^{(25+\beta_2)/2}$ and for simplicity we take $q = 2^{(25+\beta_2)/2}$. Further, following the suggestion in the second paragraph on page 7 of [19], we take $\ell = O(1)$. Then the tightness gap $G = 2^{(425+4\beta_1+21\beta_2)/2}$. Suppose we take $\beta_1 = 0$ and as in §4.7 we consider three values of β_2 , namely 128, 50 and 25. The corresponding values of the gap G are $2^{1556.5}$, $2^{737.5}$ and 2^{475} respectively. One may repeat the calculation using other values of β_1 and β_2 ; for example, taking $\beta_1 = \beta_2 = 128$, the value of G is $2^{1812.5}$.

To interpret Theorem 9 let's recall how one determines whether a security reduction from a problem \mathcal{Q} (a problem that's believed to be hard) to a problem \mathcal{P} (the problem that our cryptosystem is based on) with tightness gap G provides an assurance of security for given parameters. For us $\mathcal{Q} = K$ -SIVP $_{\gamma}$ and $\mathcal{P} = \text{ring-DLWE}_{q,\tau}$. Suppose that the fastest known algorithm to solve \mathcal{Q} with our parameters has running time T_2 . Suppose \mathcal{P} can be attacked by an algorithm taking time T_1 . Using the security reduction, we have a second algorithm for \mathcal{Q} that takes time GT_1 . We now make the reasonable assumption that this second algorithm will not set a new record for speed in solving \mathcal{Q} , and hence $GT_1 \geq T_2$, and so $T_1 \geq T_2/G$. If T_2 is exponential in the parameters and G is fairly small, we can feasibly choose the parameters so that $T_2/G \geq 2^{128}$. In this way practice-oriented provable security can provide convincing evidence of security against mathematical attacks, i.e., attacks that solve \mathcal{P} .

If we carry this out for Theorem 9 with the practical value $n = 2^{10}$, we find that the security reduction is worthless as an assurance of security. Since K -SIVP $_{\gamma}$ has not yet been extensively studied, let's take the value of T_2 from a harder problem that has been investigated at length, namely SVP. Of course, K -SVP $_{\gamma}$, which is equivalent to K -SIVP $_{\gamma}$, is presumably easier than exact-SVP for general lattices, and so we're likely to overestimate T_2 . According to [17], the fastest classical algorithms have heuristic running time $2^{0.337n+o(n)}$ and the fastest quantum algorithms have heuristic running time $2^{0.286n+o(n)}$. Taking $n = 1024$ as before and assuming that the $o(n)$ term doesn't add more than 50 to the exponent, we'll take $T_2 = 2^{395}$ for classical SVP-algorithms and $T_2 = 2^{343}$ for quantum SVP-algorithms. We won't worry about the fact that the T_2 for K -SVP $_{\gamma}$ is probably much less, and we'll use the smallest value obtained above for the gap G . With the classical and quantum values we're using for T_2 , we get the following lower bounds for the time T_1 needed to break ring-DLWE $_{q,\tau}$:

$$T_1 > 2^{-80} \quad (\text{classical}); \quad T_1 > 2^{-132} \quad (\text{quantum}). \quad (21)$$

So the theorem gives us no assurance at all.

Remark 13. *The negative exponents in the lower bounds (21) would be even more extreme if we chose a more realistic estimate for T_2 , such as $2^{n/k}$ (see Table 1) and a more cautious value for β_2 , say $\beta_2 = 128$. On the other hand, we could get a reasonable lower bound for T_1 for ring-DLWE $_{q,r}$ by increasing n , just as F. Gates [13] did for Regev’s reductions for general lattices. Because the tightness gap is so large, n would need to be significantly larger than it was for Gates. For example, if $n \approx 2^{17.5}$, then we have*

$$q > 2^{85}, \quad \gamma = 2^k > 2^{94}, \quad n/k < 1970, \quad G \approx 2^{1715},$$

1 leading to a lower bound for T_1 of approximately 2^{255} .

2 However, there are two difficulties with choosing n so large. In the first place, the efficiency advantage of ideal
3 lattices would be lost if one has to use lattices of dimension $> 185,000$. In the second place, the quantum part of
4 the reduction, which requires at least $3n^2$ logical qubits, becomes even farther removed from what can reasonably
5 be expected to be feasible. The quantum part would need 10^{11} logical qubits, roughly 20 million times as many as
6 Shor’s algorithm needs to factor a 2048-bit RSA modulus. Using a rough comparison with Shor’s algorithm (as
7 in (24) below), we estimate that the number of physical gates required would be about 2^{116} . There’s a steep price
8 to be paid for significantly increasing n .

9 5.1 Comparison to the tightness gap in the reduction from SIVP to DLWE for general 10 lattices

11 The tightness gap of the reduction from SIVP to DLWE for general lattices in [29] was analysed in [10, 33, 13].
12 All of these prior works overlooked certain aspects of the tightness gap. Taking these into consideration increases
13 the previous estimates. On the other hand, it is possible to reduce the gap by adjusting certain parameter choices
14 in the reduction in [29], which was not done in [10, 33, 13]. We first obtain a more accurate estimate of the
15 tightness gap in [29] and then compare this estimate to the tightness gap in [19]. In the description below, n is
16 the dimension of the underlying lattice and q is the modulus of the LWE problem.

17 **Tightness gap in the reduction of SIVP to search-LWE.** An algorithm \mathcal{B}_0 to solve SIVP can be constructed
18 using an algorithm \mathcal{B}_1 to solve DGS with a tightness gap of n^3 as in §3.1. An algorithm \mathcal{B}_1 to solve DGS can
19 be constructed using an algorithm \mathcal{B}_5 to solve search-LWE using algorithms \mathcal{B}_2 , \mathcal{B}_3 and \mathcal{B}_4 as intermediate
20 algorithms. The description of \mathcal{B}_1 is similar to the description of \mathcal{A}_1 given in §3.2. \mathcal{B}_1 first prepares a list of I
21 DGS samples of width large enough so that it can do so without invoking the LWE oracle. Then it goes through
22 a loop over $i_0 = 2n + \lceil (\log n)/2 \rceil$ iterations. In each iteration, it updates the list of I samples with a list of
23 another I samples of width at most half of the previous width. This is done by calling a quantum algorithm \mathcal{B}_2
24 which in turn requires the reverse of a BDD solver \mathcal{B}_3 . The BDD solver \mathcal{B}_3 is constructed using a restricted kind
25 of BDD solver \mathcal{B}_4 . This special BDD solver \mathcal{B}_4 uses the LWE solver \mathcal{B}_5 .

26 \mathcal{B}_1 calls \mathcal{B}_2 a total of $i_0 I \approx nI$ times; \mathcal{B}_2 calls the reverse of \mathcal{B}_3 once; \mathcal{B}_3 calls \mathcal{B}_4 a total of n times; and the
27 number of times \mathcal{B}_4 calls \mathcal{B}_5 is a constant multiple of I^2 . Here I is the number of LWE samples required by \mathcal{B}_5 .
28 So the total tightness gap in the reduction of SIVP to search-LWE is $n^5 I^3$.

29 **Tightness gap in the reduction of search-LWE to average-case DLWE.** Algorithm \mathcal{B}_5 to solve search-
30 LWE can be constructed using an algorithm \mathcal{B}_6 to solve worst-case DLWE with a tightness gap of qn . Both \mathcal{B}_5
31 and \mathcal{B}_6 require the same number of LWE samples. Algorithm \mathcal{B}_6 can be constructed using a (δ_1, δ_2) -distinguisher
32 \mathcal{D} to solve average-case DLWE. The tightness gap of this reduction is $I_1 I_2$ and the number of LWE samples
33 required by \mathcal{B}_6 is $I_1 I_2 \ell$, where I_1 and I_2 are constant multiples of δ_1^{-1} and δ_2^{-2} respectively, and ℓ is the number
34 of LWE samples required by \mathcal{D} . So the number I of LWE samples required by \mathcal{B}_5 is $I_1 I_2 \ell$.

35 **Overall tightness gap.** Combining the above two tightness gaps, the overall tightness gap comes out to be

$$36 \quad qn^6 \ell^3 \cdot (\delta_1 \delta_2^2)^{-4}. \tag{22}$$

1 Following the discussion in Section 5 of [29], q is a prime between n^2 and $2n^2$ and $\ell = \tilde{O}(n)$. Taking q to be
 2 about n^2 and ℓ to be about n , the expression given by (22) reduces to

$$n^{11} \cdot (\delta_1 \delta_2^2)^{-4}. \quad (23)$$

3 **Remark 14.** *The above estimate of the tightness gap improves upon the analysis in [29] in the following two*
 4 *ways.*

- 5 1. *In [29], the number of times \mathcal{B}_4 calls \mathcal{B}_5 was taken to be nI^2 so that the failure probability is at most 2^{-n} .
 6 For a concrete analysis, it is sufficient to consider the number of times \mathcal{B}_4 calls \mathcal{B}_5 to be a constant multiple
 7 of I^2 . Certainly with $n = 1024$ insisting on a 2^{-1024} failure rate would be overkill.*
- 8 2. *In [29], I_1 and I_2 were taken to be n/δ_1 and n/δ_2^2 respectively. This ensures that the failure probabilities of
 9 the worst-case to average-case reduction are asymptotically zero. The choices, however, are also an overkill.
 10 For purposes of concrete analysis it is sufficient to take I_1 and I_2 to be constant multiples of δ_1^{-1} and δ_2^{-2}
 11 respectively with constants that, for practical values of n such as 1024, are much less than n .*

12 **Remark 15.** *The two previous works [10, 13] estimated that \mathcal{B}_4 calls \mathcal{B}_5 n times and overlooked the factor I^2 ,
 13 while this factor was considered in [33]. The fact that for a concrete analysis it is sufficient to consider I^2 and not
 14 nI^2 was overlooked in [33]. All the three works [10, 33, 13] considered I to be a polynomial in n and overlooked
 15 the fact that $I = I_1 I_2 \ell$. Further, following [29] all these three works considered I_1 and I_2 to be n/δ_1 and n/δ_2^2
 16 respectively.*

17 The tightness gap in the reduction from SIVP to search-LWE for the reductions in [29] and [19] are $n^5 I^3$
 18 and $n^5 N$ respectively, where N is the number of samples required by the ring-LWE solver. Treating I and N as
 19 having similar values, the tightness gap of the reduction for general lattices is greater by a factor of about I^2 .
 20 This is due to the fact that the definition of the LWE problem in [29] requires the error to follow a fixed width
 21 Gaussian distribution. Since the width of the DGS samples is not known, \mathcal{B}_4 has to incrementally add errors so
 22 that for some increment the error distribution is negligibly far from the error distribution expected by the LWE
 23 solver. This step results in the factor I^2 arising in the tightness gap.

24 We compare the tightness gap of the reduction from approximate SIVP to DLWE for ideal lattices with
 25 that for general lattices. Taking $\delta_1 = 2^{-\beta_1}$ and $\delta_2 = 2^{-\beta_2}$, for $n = 1024$, the expression in (23) is $2^{110+4\beta_1+8\beta_2}$.
 26 For $\beta_1 = 0$ and $\beta_2 = 128, 50$ and 25 , the values of the tightness gap are 2^{1134} , 2^{510} and 2^{310} respectively. For
 27 $\beta_1 = \beta_2 = 128$, the tightness gap¹³ is 2^{1646} . The tightness gap we obtain for ideal lattices is much more than the
 28 tightness gap for general lattices.

29 **Remark 16.** *That the tightness gap for ideal lattices is even larger than for general lattices is particularly*
 30 *troubling because the gap is between the problem of cryptographic interest and a problem that is probably easier*
 31 *for ideal lattices than for general lattices.*

32 **Remark 17.** *Increasing the value of n to compensate for the tightness gap in [29] was considered in [13]. The*
 33 *above analysis shows that the estimate of the tightness gap in [13] was off the mark and so the suggested values*
 34 *of n to compensate for the gap are also off the mark. Further, as explained above, increasing the value of n to*
 35 *compensate for the tightness gap is futile, since the number of logical qubits grows quadratically in n , making an*
 36 *already infeasible quantum circuit even worse.*

37 6 Problems with the quantum part of the reduction

38 The quantum part of the reduction in [19] is largely taken from [29]. In [29] Regev writes that “This article is
 39 almost entirely classical. In fact, quantum is needed only in one step in the proof of the main theorem.” This

¹³This was the case considered in [10, 33], where the tightness gap was estimated to be 2^{524} .

1 is true from the perspective of readers who are trying to understand the description of the reduction. From a
2 pedagogical standpoint, readers who have difficulty following the construction of a quantum state — in this case
3 by applying unitary operators in a 2^N -dimensional Hilbert space with $N > 3 \times 10^6$ — can console themselves
4 with having understood the vast majority of the steps in the security reduction.

5 However, Regev’s statement could also be interpreted as suggesting that the quantum aspect has only a minor
6 effect on the feasibility of the reduction, provided that one assumes that quantum computers scaled to break RSA
7 and ECC will be possible (which is, after all, a motivation for the development of lattice-based cryptography).
8 That could not be further from the truth. In reality, the vast majority of the reduction steps must occur within
9 a quantum computer, and this raises a multitude of questions about feasibility.

The reduction is divided into a sequence of ten algorithms

$$\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_7, \mathcal{D}_1, \mathcal{D}_2.$$

10 The quantum part starts with \mathcal{A}_2 . At a key point (the first paragraph of the proof of Lemma 3.14 in [29]) the
11 quantum algorithm has a state that is a linear combination of roughly 2^{3n^2} terms, each involving two entangled
12 registers. The algorithm needs to “erase” the first of these entangled registers, which means “uncomputing” a
13 closest vector in each summand. This is done by converting an algorithm for the closest vector problem (denoted
14 CVP in [29] and BDD in [19]) into a quantum circuit and then reversing the circuit. Thus, the entire circuit
15 for the rest of the reduction from \mathcal{A}_2 to \mathcal{D}_2 has to be incorporated (after being reversed) into the quantum
16 computation. As a result, as we’ll see, the burden on the quantum computer is many orders of magnitude
17 greater than it would be for Shor’s algorithm to break RSA₂₀₄₈. Most obviously, all of the operations in the
18 sequence of eight algorithms from \mathcal{A}_2 to \mathcal{A}_7 , \mathcal{D}_1 , \mathcal{D}_2 are quantum operations, and according to [35, pp. 96-97] a
19 quantum operation can be expected to cost at least 2^{10} times as much (and possibly up to 2^{50} times as much)
20 as a classical operation.

21 6.1 The number of logical qubits

22 The quantum algorithm \mathcal{A}_2 is based on Lemma 3.14 of [29], which shows that $n \log R$ logical qubits are required
23 for an ideal \mathcal{I} , where R is an integer which is at least $2^{3n} \lambda_n(\mathcal{I})$. Since $\lambda_n(\mathcal{I})$ is generally polynomial in n , it
24 follows that the number of logical qubits required is about $3n^2$. For $n = 2^{10}$ about 3 million logical qubits will
25 be required. In comparison, factoring a 2048-bit RSA modulus requires roughly 4000 to 5000 logical qubits.

26 6.2 Circuit size

27 The size of the logical circuit — that is, neglecting error correction — depends on the length and complexity of
28 the algorithm and on the number of logical qubits of input. The algorithms from [19] surveyed in this paper form
29 a complicated interlocking sequence of computations. In addition, the number of qubits for a 1024-dimensional
30 lattice is about 750 times the number of qubits in Shor’s algorithm to break RSA₂₀₄₈. The number of gates in
31 the quantum part will be many times greater than in Shor’s algorithm, and the circuit depth will also be much
32 greater.

33 6.3 The number of physical gates

34 A basic issue in estimating the resources needed for a quantum computation is the need for error correction in
35 order to cope with quantum decoherence. This means that the number of physical gates must be many times
36 the number of logical gates. According to the definitive reference on quantum computing by M. Nielsen and
37 I. Chuang [22], much progress in quantum error correction has been made over the years, and there are reasons for
38 optimism. The main reason cited by Nielsen and Chuang for confidence that the need for error correction is not
39 an insurmountable obstacle is the “threshold theorem.” That theorem states that as long as the error probability

1 at a gate is below a certain threshold level, an arbitrarily complicated quantum circuit can be transformed into
 2 a physical circuit with negligible error, where the ratio of the number g_p of physical gates to the number g_ℓ of
 3 logical gates is polynomial in $\log(g_\ell)$.

4 We mentioned that the number of logical gates depends on the complexity of the algorithm and, for a given
 5 algorithm, on the number of qubits of input. Shor’s algorithm is simple. It consists of two computations: a
 6 modular exponentiation (or a point multiple in the case of ECC) and a quantum Fourier transform. Never-
 7 theless, the quantum resources needed to apply it to find elliptic curve discrete logarithms or factor integers of
 8 cryptographic interest are considerable. Two papers by researchers at Microsoft [31, 15] give concrete estimates
 9 for the number of qubits and the physical circuit size (number of Toffoli gates) needed to break RSA and ECC
 10 using Shor’s algorithm. Solving the Elliptic Curve Discrete Log Problem on an elliptic curve over an n -bit prime
 11 field can be done with $9n + 2\lceil \log n \rceil + 10$ qubits and $448n^3 \log n + 4090n^3$ Toffoli gates; for $n = 256$ this means
 12 2330 qubits and roughly 2^{37} Toffoli gates. Factoring an n -bit RSA modulus can be done with $2n + 2$ qubits and
 13 $64n^3 \log n + O(n^3)$ Toffoli gates; for $n = 2048$ this means 4098 qubits and roughly 1.5×2^{42} Toffoli gates. In both
 14 cases the number of Toffoli gates is roughly proportional to $n^3 \log(n)$, and since n is essentially proportional to
 15 the number of qubits, the circuit size also grows proportionally to a log term times the cube of the number of
 16 qubits.

17 Let’s try to very roughly extrapolate from integer factorization to the quantum security reductions in [29]
 18 and [19] for a 1024-dimensional lattice. Ignoring the $\log(n)$ term and also the fact that the security reductions are
 19 far more complicated than Shor’s algorithm, we can derive a very rough lower bound for the number of physical
 20 gates needed to carry out the security reduction by multiplying the number of Toffoli gates for 2048-bit integer
 21 factorization by the cube of the ratio of the number of qubits for the security reduction to the number for Shor’s
 22 algorithm:

$$(3 \cdot 1024^2 / 4098)^3 \times 1.5 \times 2^{42} \approx 2^{71} \text{ gates.} \quad (24)$$

23 Alternatively, suppose we ignore the growth of the $\text{poly}(\log(g_\ell))$ factor in the threshold theorem and make
 24 the rough assumption that the number of physical gates is simply proportional to the size of the logical circuit
 25 and that the latter is proportional to length of the algorithm times the square of the number of qubits. Further
 26 suppose that the quantum part of the security reduction is at least 25 times more complicated than Shor’s lemma.
 27 In that case we arrive at a lower bound of about $(3 \cdot 1024^2 / 4098)^2 \times 25 \times 1.5 \times 2^{42} \approx 2^{66}$ gates.

28 6.4 Internal memory

29 Besides the exorbitant cost in Toffoli gates, there’s an obstacle that’s intrinsic to the structure of the security
 30 reductions in [29] and [19]. The book by Nielsen and Chuang, after discussing the threshold theorem and
 31 describing some methods of error correction, concludes with several qualifying remarks, one of which points
 32 to the important role of interaction with a classical computer. They comment that their earlier discussion
 33 “completely neglected the cost of the classical computations and communication that are done during state
 34 preparation, syndrome measurement, and recovery. The cost of these could potentially be quite high” (p. 494).
 35 We next describe some concerns about whether it’s possible in principle for a classical computer to interact with
 36 the quantum computer during the “uncomputation” of a nearest vector in the quantum reduction.

37 Storing intermediate output while waiting for an algorithm to be ready to use it is problematic in quantum
 38 computing. Because of quantum decoherence, “No Loitering” signs are ubiquitous in the quantum computing
 39 world. Suppose that an algorithm \mathcal{A} recursively produces a sequence of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ that in a second
 40 part of the algorithm get processed in the reverse order to produce the final output. If \mathcal{A} is a standalone algorithm,
 41 it can be divided into two quantum algorithms \mathcal{A}' and \mathcal{A}'' that interact multiple times with a classical computer.
 42 The first of these runs n times. Each run outputs the next \mathbf{v}_i , which is read and stored by a classical computer
 43 that then rewires the quantum circuit using \mathbf{v}_i to prepare the circuit for computing \mathbf{v}_{i+1} . The algorithm \mathcal{A}''
 44 also runs n times, where each intermediate output is again read by the classical computer, which prepares the

1 circuit for the next run. Note that “outputting” a value to a classical computer amounts to an observation of a
2 quantum state, and that observation destroys the state.

3 This necessity of breaking up an algorithm into a sequence of sub-algorithms whose execution is terminated
4 in order to allow rewiring is very different from anything one has to worry about in classical computation, where
5 short-term storage is a basic available resource. Even Shor’s quantum algorithm, which in essence is very simple,
6 requires a rewiring for each number we want to factor.

7 Although this rewiring is key to making quantum computation work, there seems to be a fundamental obstacle
8 to such interaction with a classical computer during the course of the quantum part of the security reduction
9 in [19], because the entire quantum computation (actually, quantum “uncomputation”) is being applied to each
10 summand in the quantum state constructed by \mathcal{A}_2 . This means that it’s unclear how the recursive steps could
11 be carried out in practice.

12 The first place where this concern arises is in the reduction from \mathcal{A}_3 to \mathcal{A}_4 . Starting with the input y to \mathcal{A}_3 ,
13 that reduction recursively generates a sequence of b_i that give a sequence of vectors that are closer and closer to
14 the lattice. When the nearest lattice vector can finally be determined by Babai’s algorithm, the second part of
15 the algorithm uses that lattice vector along with the b_i ’s in reverse order to compute the closest lattice vector to
16 y . Thus, the reduction has the same structure as the algorithm \mathcal{A} at the beginning of this subsection, with the
17 crucial difference that the algorithm cannot be carried out after the first interaction with a classical computer
18 collapses the state that the algorithm is being applied to.

19 The internal storage issue again arises when we consider the reduction from \mathcal{A}_4 to \mathcal{A}_5 . In that reduction \mathcal{A}_4
20 runs its subroutine (oracle) that produces vectors from the distribution $D_{I,r}$. Then each such vector is used,
21 along with the q -BDD $_{\mathcal{I},\xi}$ -input vector y , to create a sample from $A_{s,r}$ that will be used by \mathcal{A}_5 to find s , and
22 hence y . The algorithm \mathcal{A}_5 presumably needs to have a lot of samples — that is, a lot of approximate equations
23 — before it can get to work finding s . Where are the first samples from $A_{s,r}$ kept while the later ones are being
24 generated?

25 The same problem arises in the reduction of \mathcal{A}_5 to \mathcal{A}_7 (via \mathcal{A}_6), which entails finding each of the n components
26 of s modulo q_i , which are then combined by the Chinese Remainder Theorem to find s modulo q . Each component
27 is determined by running the ring-VWDLWE $_{q,\leq\alpha}^i$ solver for up to q different possible values of the residue. The
28 Chinese Remainder Theorem has to wait for all the components to be computed. How are the first components
29 that are computed going to be preserved while the later ones are being determined?

30 Finally, in the reduction of \mathcal{A}_7 to \mathcal{D}_1 , the former calls ℓ times upon its oracle for $A_{s,r}^j$ and its oracle for $D_{\mathfrak{t}}$
31 in order to produce two lists of samples \mathcal{T}' and \mathcal{T}'' , which \mathcal{D}_1 then has to distinguish between (in the case when
32 $j = i-$). It’s reasonable to assume that \mathcal{D}_1 needs to have the first samples that were created still available after
33 receiving the rest of each list. Again we do not see a way to store that data until it is needed.

34 6.5 Contrast between classical and quantum reductions

35 In a classical reduction from a problem \mathcal{Q} to a problem \mathcal{P} , the main issue of feasibility is the tightness gap, in
36 other words, the running time of the reduction algorithm that solves \mathcal{Q} given an oracle for \mathcal{P} that runs in unit
37 time. Once the tightness gap is computed and one has a reasonable running time estimate for the best available
38 algorithm for the supposedly hard problem \mathcal{Q} , one can give key length recommendations based on the guarantees
39 that the reduction gives. In this way one realizes “practice-oriented provable security” [3].

40 In quantum reductions, on the other hand, it is not always possible to obtain a meaningful security guarantee
41 simply by increasing key length. One must also consider potential obstacles to feasibility that do not arise in
42 the classical case. As we’ve seen, one such obstacle is the size of the reduction algorithm’s physical circuit,
43 which depends on the number of qubits, the complexity of the algorithm, the nature of the tasks performed, and
44 the way error correction is implemented. In addition, the problem of internal storage requires interventions to
45 prepare the circuit for the next stage of the computations. In some cases it is unclear that this is theoretically
46 possible, let alone achievable in practice.

1 The quantum reduction for lattice-based cryptography, first described in [29], was a remarkable achievement,
2 opening up a new direction in provable security research. At the same time it also opened up new challenges in
3 determining whether or not a security reduction gives a meaningful guarantee.

4 6.6 Some caveats

5 Our discussion of the quantum part of the SIVP_γ -to-DLWE reduction is of necessity informal, imprecise, and
6 speculative. None of us are experts in quantum computation. In particular, we’re making the assumption that
7 our analysis of a quantum BDD algorithm with a DLWE-oracle also applies to the reversed circuit — that is,
8 we’re assuming that if the BDD algorithm is infeasible because of quantum decoherence, then so is the reverse
9 circuit.

10 Moreover, even the experts cannot accurately predict how successful physicists and engineers will be in the
11 coming decades in overcoming the formidable obstacles to the development of a large-scale quantum computer.
12 But in any case, after considering the issues of qubit and circuit size, error correction, and internal memory, it
13 should be clear that, even if it is possible in principle to carry out the reduction, the level of quantum scaling
14 required to do so is far, far greater than the amount needed to break RSA and ECC.

15 6.7 Summary

16 There are two reasons to doubt the feasibility of the quantum part of the security reduction in [19]. Even for
17 $n = 1024$ the circuit size is many thousands times the circuit size for Shor’s algorithm to factor a 2048-bit RSA
18 modulus. Since the number of qubits grows quadratically with n , the circuit size becomes much greater if one
19 chooses n large enough to compensate for the tightness gap in the reduction. In addition, the quantum algorithm
20 requires almost continual interaction with a classical computer, and we see no way to do this without destroying
21 the quantum state that the algorithm is being applied to.

22 **Remark 18.** *Most of the discussion in this section also applies to the quantum part of [29], upon which the*
23 *quantum part of [19] is based. For example, both quantum algorithms require at least $3n^2$ logical qubits. The*
24 *concrete analyses of Regev’s reduction in [10, 33, 13] did not look at feasibility of the quantum part.*

25 7 Subsequent work

26 We consider two works which extend the results in [19]. The first generalises from ideal lattices to module lattices
27 while the second generalises to any number field.

28 7.1 Module lattices

29 Let K be a number field of degree n and R be its ring of integers. For a positive integer d , K^d is a module
30 over R . The embedding σ of K into H extends to an embedding of K^d into H^d , and one may consider lattices
31 in H^d . The dimension of this lattice is nd . As pointed out in [18], module-LWE serves as a “bridge” between
32 ring-LWE and LWE for general lattices. If $d = 1$ we have ring-LWE for an n -dimensional ideal lattice, and if
33 $n = 1$ we have LWE for a d -dimensional general lattice. For reasons of efficiency, in practice d is a small constant;
34 in Kyber [7] and SABER [12] n is taken to be 256 and d is either 2, 3, or 4. Thus, in practice module-LWE is
35 much closer to ring-LWE than to general lattice LWE. Below we will see that the tightness gap for the reduction
36 from approximate module-SIVP to module-DLWE is a little less than for the K -SIVP $_\gamma$ to ring-DLWE reduction,
37 but considerably more than for the general SIVP $_\gamma$ to DLWE reduction. This is not surprising.

38 In [18] it is suggested that module-DLWE may be harder than ring-DLWE. This is a reasonable conjecture.
39 Indeed, if ring-DLWE is much easier than general DLWE, then the difference in hardness between module-DLWE
40 with $d = 2, 3$, or 4 and general DLWE will probably turn out to be a little less.

1 The module-LWE distribution is the following. Let $\mathbb{T} = H/\sigma(R^\vee)$ as before and $q \geq 2$ be an integer. For
2 $\mathbf{s} \in (R_q^\vee)^d$ and $\mathbf{a} \in R_q^d$, a sample from the module-LWE distribution with parameters q , \mathbf{s} and r is $(\mathbf{a}, \sigma(\langle \mathbf{a}, \mathbf{s} \rangle / q) + \mathbf{e}$
3 $\text{mod } \sigma(R_q^\vee))$, where \mathbf{e} is chosen from H following the distribution D_r . The search and decision versions of the
4 module-LWE problem are defined in a manner analogously to those of ring-LWE problem.

5 Let K be the m -th cyclotomic field having degree n , $d \geq 1$ be an integer, and q be a prime greater than
6 2 such that $q \equiv 1 \pmod{m}$. The reduction in [19] from approximate ideal-SIVP to ring-DLWE was generalised
7 to module lattices in [18], where it was shown that approximate module-SIVP reduces to module-DLWE. The
8 sequence of algorithms in the reduction in [18] is exactly the same as in the reduction in [19]. In this section
9 we let \mathcal{A}_0 to \mathcal{A}_7 and \mathcal{D}_1 and \mathcal{D}_2 denote the algorithms for module-SIVP to module-DLWE reduction. Here \mathcal{A}_0
10 solves module-SIVP $_\gamma$ and \mathcal{D}_2 is a (δ_1, δ_2) -distinguisher for the module-DLWE problem with parameters q and \mathbf{r} .
11 Below we briefly summarise the parameters and the tightness gap of the reduction.

12 The number of times \mathcal{A}_0 calls \mathcal{A}_5 is about $(nd)^5 \cdot N$, where N is the number of module-LWE samples required
13 by \mathcal{A}_5 . The number of times \mathcal{A}_5 calls \mathcal{A}_7 is qnd and the number of module-LWE samples required by \mathcal{A}_7 and \mathcal{A}_5
14 are equal. The number of times \mathcal{A}_7 calls the (ϵ_1, ϵ_2) -distinguisher \mathcal{D}_1 is $N_1 N_2$ and the number of module-LWE
15 samples required by \mathcal{A}_7 is $N_1 N_2 \ell$, where $N_1 = (nN_2 \ell)^{1/2} / (\epsilon_1 \epsilon_2^2)$, $N_2 = \tilde{O}(\epsilon_2^{-2})$ and ℓ is the number of LWE-
16 samples required by \mathcal{D}_1 . So $N = N_1 N_2 \ell$. The (ϵ_1, ϵ_2) -distinguisher \mathcal{D}_1 can be used as a (δ_1, δ_2) distinguisher \mathcal{D}_2 ,
17 where $\epsilon_1 = \delta_1/n$ and $\epsilon_2 = \delta_2/n$. The entire reduction from module-SIVP to module-DLWE needs to be invoked
18 an additional n times due to the reason noted in Remark 11. The overall tightness gap is about

$$qn^{20} d^6 \ell^2 (\delta_1 \delta_2^5)^{-2}. \quad (25)$$

19 Also, the quantum part of the reduction requires about $(nd)^2$ logical qubits.

20 For fixed dimension nd the tightness gap in (25) for the module case is less than that in (20) for the ring case
21 by a factor of d^{14} . Thus, for the values $n = 256$ and $d = 4$ – the values given for a high security level for Kyber
22 and SABER – the lower bounds for T_1 in (21) should be replaced by

$$T_1 > 2^{-52} \text{ (classical); } \quad T_1 > 2^{-104} \text{ (quantum).}$$

23 **Remark 19.** *The discussion in §6 regarding the quantum part of the reduction of approximate ideal-SIVP to*
24 *ring-DLWE in [19] also applies to the reduction of approximate module-SIVP to module-DLWE in [18].*

25 7.2 Arbitrary number fields

26 A follow-up work [27] (whose latest version is [26]) improved upon the reduction in [19]. The reduction of
27 K -SIVP $_\gamma$ to ring-DLWE $_{q,r}$ in [19] holds only for cyclotomic number fields. In [27, 26] this reduction was extended
28 to any number field. Unlike [19], the reduction in [27, 26] does not go through the search ring-LWE problem.
29 Below we discuss the tightness gap in the reduction in [27, 26]. Specifically, we refer to the version in [26].

30 As in [19], the reduction of K -SIVP $_\gamma$ to ring-DLWE $_{q,r}$ in [26] also goes through several steps. The first step
31 is to reduce K -SIVP $_\gamma$ to K -DGS $_\Gamma$ and is exactly the algorithm \mathcal{A}_0 described earlier that has a tightness gap
32 of n^3 . The reduction of K -DGS $_\Gamma$ to ring-DLWE $_{q,r}$ is given by a sequence of algorithms. The first of these is
33 the algorithm \mathcal{A}_1 described earlier. Briefly, \mathcal{A}_1 prepares an initial list of N DGS samples and goes through
34 $i_0 = (2n + (\log n)/2)$ iterations, where in each iteration \mathcal{A}_1 invokes the quantum circuit \mathcal{A}_2 N times, and in each
35 invocation \mathcal{A}_1 provides \mathcal{A}_2 with a list of DGS samples and receives in return DGS samples with width reduced
36 by a factor of at least 2. Finally, \mathcal{A}_1 returns a sample from the last list that it prepares. The number N of DGS
37 samples is equal to the number of LWE samples required by a distinguisher for the ring-DLWE $_{q,r}$ problem, and
38 this number in [26] is very different from that of [19]. As before, \mathcal{A}_1 calls \mathcal{A}_2 a total of $(2n + (\log n)/2)N$ times.

39 From this point onwards, the reductions in [19] and [26] begin to differ. In [26], \mathcal{A}_2 applies the reverse of an
40 algorithm \mathcal{B}_3 that solves the Gaussian decoding problem¹⁴. The construction of \mathcal{B}_3 is based on a distinguisher

¹⁴For a lattice $\Lambda \subset H$ and a parameter ξ , an instance of the problem is a coset $\mathbf{e} + \Lambda$, where \mathbf{e} is drawn from D_ξ , and the task is to find \mathbf{e} . Note that this is essentially our definition of the BDD problem in Section 3.

1 \mathcal{E}_1 for the ring-DLWE $_{q,\leq\alpha}$ problem¹⁵. The construction of \mathcal{E}_1 is based on a (δ_1, δ_2) -distinguisher \mathcal{D}_2 for the ring-
2 DLWE $_{q,\tau}$ problem where the relation between τ and α is given by (11). The construction of \mathcal{E}_1 from \mathcal{D}_2 is very
3 similar to the construction of \mathcal{A}_7 from \mathcal{D}_1 described in §4.4, and the concreteness aspects are exactly the same.
4 In particular, \mathcal{E}_1 calls \mathcal{D}_2 about $(\delta_1\delta_2^5)^{-1}(n\ell)^{1/2}$ times and the number M_1 of LWE samples required by \mathcal{E}_1 is
5 about $(\delta_1\delta_2^5)^{-1} \cdot (n^{1/2}\ell^{3/2})$, where ℓ is the number of LWE samples required by \mathcal{D}_2 .

6 The main technical contribution of [26] is the construction of \mathcal{B}_3 from \mathcal{E}_1 . We will not get into the details of
7 this very complicated construction, and instead we simply identify the tightness gap of this reduction. Let M_2
8 be the number of times \mathcal{B}_3 calls \mathcal{E}_1 . From the description in [26], we obtain an estimate of M_2 . Note that each
9 call to \mathcal{E}_1 requires M_1 samples so that the total number of samples required in all the M_2 calls is equal to M_1M_2 .
10 This is the value of N , since M_1M_2 DGS samples need to be provided to \mathcal{B}_3 so that it can generate the required
11 number of LWE samples.

12 Let¹⁶ $\kappa = \text{poly}(n) \geq 100n^2M_1$ and $\mu = \text{poly}(\kappa)$. Recall that $n = s_1 + 2s_2$, where s_1 and $2s_2$ are the numbers
13 of real and complex roots respectively of the defining polynomial of the number field. Using \mathcal{E}_1 , \mathcal{B}_3 creates $s_1 + s_2$
14 oracles, where each of these oracles calls \mathcal{E}_1 once (Lemma 6.6 in [26]). Corresponding to these $s_1 + s_2$ oracles,
15 $s_1 + s_2$ algorithms are created, where (ignoring logarithmic factors) each algorithm calls its corresponding oracle
16 about $5 \times 10^{15} \cdot \kappa^6 \mu^3$ times, where κ and μ are at least $100n^2M_1$ (first part of the proof of Proposition 4.4 and
17 the proof of Lemma 6.6 of [26]). Each of these $s_1 + s_2$ algorithms is itself called about $2000\kappa^3$ times (second part
18 of the proof of Proposition 4.4 of [26]) For simplicity, we take $\mu = \kappa = 100n^2M_1$ and take $2(s_1 + s_2)$ to be n to
19 obtain the value of M_2 to be about $10^{43} \cdot n^{25}M_1^{12}$.

20 The overall tightness gap is given by the number of times \mathcal{A}_0 calls \mathcal{D}_2 . This number is about

$$\begin{aligned}
n^3 \cdot (2n + (\log n)/2)N \cdot M_2 \cdot (\delta_1\delta_2^5)^{-1}(n\ell)^{1/2} &\approx n^4 \cdot M_1M_2^2 \cdot (\delta_1\delta_2^5)^{-1}(n\ell)^{1/2} \\
&\approx 10^{86} \cdot n^{54} \cdot M_1^{25} \cdot (\delta_1\delta_2^5)^{-1}(n\ell)^{1/2} \\
&\approx 10^{86} \cdot n^{54} \cdot ((\delta_1\delta_2^5)^{-1} \cdot (n^{1/2}\ell^{3/2}))^{25} \cdot (\delta_1\delta_2^5)^{-1}(n\ell)^{1/2} \\
&= 10^{86} \cdot n^{67} \cdot \ell^{38} \cdot (\delta_1\delta_2^5)^{-26}.
\end{aligned} \tag{26}$$

21 **Remark 20.** *The discussion in §6 regarding the quantum circuit required for the reduction in [19] applies equally*
22 *to the reduction in [26]*

23 The estimate of the tightness gap given by (26) shows that from a practical point of view the reduction
24 is completely meaningless. This estimate may also be compared with the estimate of the tightness gap of the
25 reduction in [19] given by (20) and the tightness gap of the reduction in [29] given in [10, 33] (see the discussion
26 following Theorem 9). While the tightness gap of the reduction in [29] is itself huge, it is lower than the tightness
27 gap of the reduction in [19], and the tightness gap of the reduction in [26] is much much larger than the tightness
28 gaps of the reductions in both [29] and [19].

29 **Remark 21.** *It was remarked in [26] (sentence before Definition 4.1) that the authors did not try to optimise*
30 *the parameters. So it is possible that the estimate given by (26) can be lowered with optimised parameters.*

31 8 Conclusion

32 Our main result is that the security reduction in [19] gives no meaningful guarantee of real-world security for three
33 reasons – a huge tightness gap, obstacles to realisation of the quantum part, and an approximation factor for
34 the K -SIVP $_\gamma$ problem that makes it unlikely that the problem is hard, particularly since the lattice problem has
35 been restricted to a subset of lattices with special geometric and algebraic properties. In addition, the reduction
36 requires a value for the modulus q that is much larger than in proposed implementations.

¹⁵Formally, it is necessary to consider elliptical Gaussian distribution, but this does not matter for the tightness analysis.

¹⁶The constant 100 is from [26].

1 But we need to qualify this by clarifying what we are not claiming. First, we do not have an actual attack
2 on LWE security, and are not saying that one necessarily exists. Nor are we saying that the tightness gap in the
3 reduction in [19] is intrinsic to the approximate ideal-SIVP $_{\gamma}$ and DLWE problems; it is certainly possible that a
4 reduction could be constructed with a lower tightness gap. Similarly, we invite the reader to look for a reduction
5 of ring-LWE to ring-DLWE with \mathfrak{r}/α equal to $(n\ell/\ln(n\ell))^{1/4}$ as claimed in [19] rather than $(nN_2\ell/\ln(nN_2\ell))^{1/4}$,
6 which was the best we could rigorously justify.

7 In theory, it would be possible to compensate for the tightness gap by increasing n . However, this is not so
8 simple as for classical reductions because, as remarked in Section 5, it entails a sharp increase in the number of
9 logical qubits and physical gates in the quantum part of the reduction. Although the number of logical qubits is
10 “only” quadratic in n , in the quantum context that rate of growth gives rise to grave doubts about feasibility.

11 We have no intention of questioning the quality of the work in [19]. The construction of the security reduction
12 involving ten nested sub-algorithms was a true *tour de force*. From a theoretical standpoint, it was a major
13 accomplishment to construct a polynomial time reduction from approximate ideal-SIVP $_{\gamma}$ to the decision problem
14 that proposed cryptosystems are based on. Unfortunately, the reduction loses its value when viewed from the
15 vantage point of practice-oriented provable security.

16 The four major problems in the reduction in [19] (the large tightness gap, the large value of the approximation
17 factor, the unrealistic quantum part, and the likelihood that approximate ideal-SIVP is substantially easier than
18 approximate SIVP for general lattices) mean that one cannot have any confidence that the reduction rules
19 out practical mathematical attacks on ring-DLWE based cryptosystems. It seems that similar doubts apply to
20 module-DLWE based systems, although it is reasonable to conjecture that such a system with $n = 256$ and $d = 4$
21 is a little less vulnerable than a ring-DLWE based system with $n = 1024$.

22 In 2015 Peikert¹⁷ discussed asymptotic analyses of the security of lattice-based systems and concluded that
23 they ensure the superiority of such systems from the standpoint of security:

24 ...worst-case reductions give a hard-and-fast guarantee that the cryptosystem is at least as hard to
25 break as the hardest instances of some underlying problem. This gives a true lower bound on security,
26 and prevents the kind of unexpected weaknesses that have so often been exposed in schemes that lack
27 such reductions.

28 But the type of security reductions analysed in this paper give no guarantees of real-world security, let alone
29 “hard-and-fast” ones. There is no meaningful “true lower bound” on security. Among the proposals for post-
30 quantum cryptography, the ones that have badly deficient “proofs of security” should not be privileged over those
31 that have arguments for security that are grounded in heuristics and practical analysis.

32 Acknowledgements

33 We wish to thank Dan Bernstein and Alfred Menezes for helpful comments on an earlier draft and Ann Hibner
34 Kobitz for editorial corrections and comments. Of course, any opinions or errors in the paper are the sole
35 responsibility of the authors.

36 References

- 37 [1] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - a new
38 hope. Cryptology ePrint Archive, Report 2015/1092, 2015. <https://ia.cr/2015/1092>.
- 39 [2] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische*
40 *Annalen*, 296:625–636, 1993.

¹⁷C. Peikert, 19 February 2015 blog posting, <http://web.eecs.umich.edu/~cpeikert/soliloquy.html>; accessed on March 5, 2022.

- 1 [3] Mihir Bellare. Practice-oriented provable-security. In *International Workshop on Information Security*,
2 volume 1396 of *Lecture Notes in Computer Science*, pages 221–231. Springer, 1997.
- 3 [4] Daniel J. Bernstein. Comparing proofs of security for lattice-based encryption. Cryptology ePrint Archive,
4 Report 2019/691, 2019. <https://eprint.iacr.org/2019/691>.
- 5 [5] Daniel J. Bernstein and Tanja Lange. Non-randomness of S-unit lattices. Cryptology ePrint Archive, Report
6 2021/1428, 2021. <https://eprint.iacr.org/2021/1428>.
- 7 [6] Joppe W. Bos, Craig Costello, Michael Naehrig, and Douglas Stebila. Post-quantum key exchange for the
8 TLS protocol from the ring learning with errors problem. In *IEEE Symposium on Security and Privacy*,
9 pages 553–570. IEEE Computer Society, 2015.
- 10 [7] Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter
11 Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - kyber: A cca-secure module-lattice-based KEM.
12 In *EuroS&P*, pages 353–367. IEEE, 2018.
- 13 [8] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of
14 learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on
15 Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 575–584. ACM,
16 2013.
- 17 [9] George Casella and Roger L. Berger. *Statistical Inference*. Thomson, 2002.
- 18 [10] Sanjit Chatterjee, Neal Koblitz, Alfred Menezes, and Palash Sarkar. Another look at tightness II: prac-
19 tical issues in cryptography. In Raphael C.-W. Phan and Moti Yung, editors, *Paradigms in Cryptology
20 - Mycrypt 2016. Malicious and Exploratory Cryptology - Second International Conference, Mycrypt 2016,
21 Kuala Lumpur, Malaysia, December 1-2, 2016, Revised Selected Papers*, volume 10311 of *Lecture Notes in
22 Computer Science*, pages 21–55. Springer, 2016.
- 23 [11] Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Mildly short vectors in cyclotomic ideal lattices in
24 quantum polynomial time. *J. ACM*, 68(2):8:1–8:26, 2021.
- 25 [12] Jan-Pieter D’Anvers, Angshuman Karmakar, Sujoy Sinha Roy, and Frederik Vercauteren. Saber: Module-
26 lwr based key exchange, cpa-secure encryption and cca-secure KEM. In *AFRICACRYPT*, volume 10831 of
27 *Lecture Notes in Computer Science*, pages 282–305. Springer, 2018.
- 28 [13] Fletcher Gates. Reduction-respecting parameters for lattice-based cryptosystems. [https://macsphere.
29 mcmaster.ca/bitstream/11375/24466/2/gates_fletcher_m_finalsubmission2018october_msc.pdf](https://macsphere.mcmaster.ca/bitstream/11375/24466/2/gates_fletcher_m_finalsubmission2018october_msc.pdf),
30 2018. Masters Thesis, McMaster University.
- 31 [14] Oded Goldreich and Shafi Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput.
32 Syst. Sci.*, 60(3):540–563, 2000.
- 33 [15] Thomas Häner, Martin Roetteler, and Krysta M. Svore. Factoring using $2n + 2$ qubits with Toffoli based
34 modular multiplication. arXiv preprint <https://arxiv.org/abs/1611.07995>, 2016.
- 35 [16] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American
36 Statistical Association*, 58(301):13–30, 1963.
- 37 [17] Thijs Laarhoven, Michele Mosca, and Joop van de Pol. Finding shortest lattice vectors faster using quantum
38 search. *Designs, Codes and Cryptography*, 77(2-3):375–400, 2015.

- 1 [18] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Designs,*
2 *Codes and Cryptography*, 75(3):565–599, 2015.
- 3 [19] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings.
4 *J. ACM*, 60(6):43:1–43:35, 2013.
- 5 [20] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: A Cryptographic Perspective*. Springer,
6 2002.
- 7 [21] Daniele Micciancio. Lattice algorithms and applications: Basis reduction. [https://cseweb.ucsd.edu/
8 classes/sp14/cse206A-a/lec5.pdf](https://cseweb.ucsd.edu/classes/sp14/cse206A-a/lec5.pdf), 2014. Accessed on February 8, 2022.
- 9 [22] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge
10 University Press, 2nd edition, 2010.
- 11 [23] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In
12 Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing,*
13 *STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 333–342. ACM, 2009.
- 14 [24] Chris Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*,
15 10(4):283–424, 2016. <https://doi.org/10.1561/04000000074>.
- 16 [25] Chris Peikert and Zachary Pepin. Algebraically structured LWE, revisited. In *TCC (1)*, volume 11891 of
17 *Lecture Notes in Computer Science*, pages 1–23. Springer, 2019.
- 18 [26] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring
19 and modulus. *IACR Cryptol. ePrint Arch.*, 2017. <https://eprint.iacr.org/2017/258>, version dated 6
20 June, 2020.
- 21 [27] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring
22 and modulus. In *STOC*, pages 461–473. ACM, 2017.
- 23 [28] Oded Regev. Lattices in computer science: Average-case hardness. [https://cims.nyu.edu/~regev/
24 teaching/lattices_fall_2004/ln/averagecase.pdf](https://cims.nyu.edu/~regev/teaching/lattices_fall_2004/ln/averagecase.pdf), accessed on February 8, 2022, 2004.
- 25 [29] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–
26 34:40, 2009.
- 27 [30] Oded Regev. The learning with errors problem (invited survey). In *Computational Complexity Conference*,
28 pages 191–204. IEEE Computer Society, 2010.
- 29 [31] Martin Roetteler, Michael Naehrig, Krysta M. Svore, and Kristin E. Lauter. Quantum resource estimates
30 for computing elliptic curve discrete logarithms. In *ASIACRYPT (2)*, volume 10625 of *Lecture Notes in*
31 *Computer Science*, pages 241–270. Springer, 2017.
- 32 [32] Palash Sarkar and Subhadip Singha. Classical reduction of gap SVP to LWE: A concrete security analysis.
33 *Advances in Mathematics of Communications*, 2021. [https://www.aimsciences.org/article/doi/10.
34 3934/amc.2021004](https://www.aimsciences.org/article/doi/10.3934/amc.2021004).
- 35 [33] Palash Sarkar and Subhadip Singha. Verifying solutions to LWE with implications for concrete security.
36 *Advances in Mathematics of Communications*, 15(2):257–266, 2021.

[34] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, volume 5912 of *Lecture Notes in Computer Science*, pages 617–635. Springer, 2009.

[35] NTRU Prime Risk-Management Team. Risks of lattice KEMs. 2021. <https://ntruprime.cr.yp.to/latticerisks-20211031.pdf>, accessed on February 18, 2022.

A Details of the parameters of \mathcal{A}_2 and \mathcal{A}_3

Algorithm \mathcal{A}_2 takes as input an ideal \mathcal{I} and a set of samples from $D_{\mathcal{I},r}$ and returns a sample from $D_{\mathcal{I},r'}$. Algorithm \mathcal{A}_3 takes as input a pair (\mathcal{I}^\vee, y) , where \mathcal{I} is a fractional ideal of K , $y = x + e$, $x \in \mathcal{I}^\vee$ and $e = \sigma^{-1}(\mathbf{e})$ with \mathbf{e} being chosen according to the distribution D_ξ . Additionally, \mathcal{A}_3 also has access to a set of samples from $D_{\mathcal{I},r}$. The parameters¹⁸ r, r' and ξ have to satisfy the following relations.

$$\left. \begin{aligned} r &\geq \sqrt{2}q \cdot \eta_\varepsilon(\mathcal{I}), \\ r' &= r \cdot \omega(\sqrt{\ln n}) / (\alpha q) > \sqrt{2n} / \lambda_1(\mathcal{I}^\vee), \\ \xi &= (\alpha q) / (2r \cdot \omega(\sqrt{\ln n})) < \lambda_1(\mathcal{I}^\vee) / (2\sqrt{2}n). \end{aligned} \right\} \quad (27)$$

Let $\xi' = \xi \cdot \sqrt{2} \cdot \omega(\sqrt{\ln n}) = (\alpha q) / (\sqrt{2}r)$. For \mathbf{e} chosen according to D_ξ , $\|\mathbf{e}\|_\infty \leq \xi'$ except with negligible probability.

The input to Algorithm \mathcal{A}_1 is a pair (\mathcal{I}, r) where $r \geq \Gamma(\mathcal{I})$ and $\Gamma(\mathcal{I})$ is given by (6). Recall that $i_0 = 2n + \lceil (\log n)/2 \rceil$ and $r_i = r \cdot (\alpha q / \omega(\sqrt{\ln n}))^i \geq 2^i r$ for $i = 0, \dots, i_0$ and $\xi_i = (\alpha q) / (2r_i \cdot \omega(\sqrt{\ln n}))$ for $i = 1, \dots, i_0$. Also, we have $\xi'_i = \alpha q / (\sqrt{2}r_i)$. Note that $r_{i-1} = r_i \cdot \omega(\sqrt{\ln n}) / (\alpha q) \leq r_i/2$ for $i = 1, \dots, i_0$, and \mathcal{S}_i is a list of independent samples from $D_{\mathcal{I},r_i}$, for $i = 0, \dots, i_0$.

Algorithm \mathcal{A}_1 first prepares \mathcal{S}_{i_0} . Then for each i going down from i_0 to 1, \mathcal{A}_1 gives \mathcal{A}_2 the input \mathcal{I} and the set \mathcal{S}_i and receives in return a sample from $D_{\mathcal{I},r_{i-1}}$. This is done N times to prepare the list \mathcal{S}_{i-1} . Since the samples $\mathcal{S}_{i_0}, \dots, \mathcal{S}_1$ are provided as input to \mathcal{A}_2 , the values r_{i_0}, \dots, r_1 have to satisfy the condition on r given in (27). Further, since \mathcal{A}_2 provides samples from $D_{\mathcal{I},r_{i_0-1}}, \dots, D_{\mathcal{I},0}$ as output, the values r_{i_0-1}, \dots, r_0 have to satisfy the condition on r' given in (27). Algorithm \mathcal{A}_3 is invoked on instances $(\mathcal{I}^\vee, y_{i_0}), \dots, (\mathcal{I}^\vee, y_1)$, where the offsets e_{i_0}, \dots, e_1 in the y_{i_0}, \dots, y_1 are sampled respectively from $D_{\xi_{i_0}}, \dots, D_{\xi_1}$ so that the values ξ_{i_0}, \dots, ξ_1 have to satisfy the conditions on ξ given in (27). The conditions on r_i and ξ_i are shown below. Before that we show $r_{i_0} \geq 2^{2n} \lambda_n(\mathcal{I})$, so that \mathcal{A}_1 can prepare \mathcal{S}_{i_0} directly (i.e., without invoking \mathcal{A}_2).

Claim 2.13 of [29] shows that the inequality $\eta_\varepsilon(\mathcal{I}) > 1/\lambda_1(\mathcal{I}^\vee)$ holds for $\varepsilon \leq e^{-\pi}$. Under the assumptions $\varepsilon \leq e^{-\pi}$ and $\alpha < \sqrt{\ln n/n}$, it follows that

$$\Gamma(\mathcal{I}) > \sqrt{2n} / \lambda_1(\mathcal{I}^\vee). \quad (28)$$

The following computation shows that $r_{i_0} \geq 2^{2n} \lambda_n(\mathcal{I})$.

$$\begin{aligned} r_{i_0} &\geq 2^{i_0} \cdot r \\ &\geq 2^{2n + \frac{1}{2} \log n} \cdot \frac{\sqrt{2} \cdot \omega(\sqrt{\ln n}) \cdot \eta_\varepsilon(\mathcal{I})}{\alpha} \quad (\text{using (6)}) \\ &\geq \sqrt{n} 2^{2n} \cdot \frac{\sqrt{2} \cdot \omega(\sqrt{\ln n})}{\alpha} \cdot \sqrt{\frac{\ln 1/\varepsilon}{\pi}} \cdot \frac{\lambda_n(\mathcal{I})}{n} \quad (\text{using Claim 2.13 of [29]}^{19}) \\ &> 2^{2n} \lambda_n(\mathcal{I}) \quad (\text{using } \omega(\sqrt{\ln n}) > \sqrt{\ln n}, \alpha < \sqrt{\ln n/n} \text{ and } \varepsilon < e^{-\pi}). \end{aligned}$$

For $i = 1, \dots, i_0$, we have the following.

¹⁸In terms of the notation used in Lemma 4.3, 4.4 and 4.7 of [19], $d' = \xi \cdot \sqrt{2n}$ and $d = \xi'$.

¹⁹Claim 2.13 of [29] shows that $\eta_\varepsilon(\mathcal{I}) \geq \sqrt{\ln(1/\varepsilon)/\pi} \cdot \lambda_n(\mathcal{I})/n$.

- 1 • Since $r \geq \Gamma(\mathcal{I})$, using the definition of $\Gamma(\mathcal{I})$ in (6), we have $r_i \geq r \cdot (\alpha q)/\omega(\sqrt{\ln n}) \geq \sqrt{2}q \cdot \eta_\epsilon(\mathcal{I})$.
- 2 • Since $r_{i-1} \geq r$, and $r \geq \Gamma(\mathcal{I})$, using (28), we have $r_{i-1} > \sqrt{2n}/\lambda_1(\mathcal{I}^\vee)$.
- 3 • Noting that $\xi_i = 1/(2r_{i-1})$, using the previous point, we have $\xi_i < \lambda_1(\mathcal{I}^\vee)/(2\sqrt{2n})$.

4 B Details of the analysis in Section 4.4

5 We first show the lower bound on the probability of a good \mathbf{z} for a good $s + t$.

6 **Proposition 10.** *For a good $s + t$, under the error distribution $D_{\tau}^{\ell N_2}$ the probability that \mathbf{z} is good is at least*
 7 $\epsilon_2/4$.

8 *Proof.* Since $s + t$ is good, from (14) we have $|\hat{p}_{s+t,0} - \hat{p}_{s+t,1}| \geq \epsilon_2/2$. Without loss of generality, we assume
 9 $\hat{p}_{s+t,0} \geq \hat{p}_{s+t,1} + \epsilon_2/2$. Let Z denote the set of all $N_2\ell$ -tuples \mathbf{z} , and let Y denote the subset of Z consisting of \mathbf{z}
 10 such that $\hat{p}_{s+t,\mathbf{z},0} \geq \hat{p}_{s+t,\mathbf{z},1} + \epsilon_2/4$. We claim that Y has measure at least $\epsilon_2/4$. Assume the contrary. We then
 11 have

$$\begin{aligned}
 \hat{p}_{s+t,0} &= \int_Y \hat{p}_{s+t,\mathbf{z},0} D_{\tau}^{\ell N_2}(\mathbf{z}) + \int_{Z \setminus Y} \hat{p}_{s+t,\mathbf{z},0} D_{\tau}^{\ell N_2}(\mathbf{z}) \\
 &< \int_Y 1 \cdot D_{\tau}^{\ell N_2}(\mathbf{z}) + \int_Z (\hat{p}_{s+t,\mathbf{z},1} + \epsilon_2/4) D_{\tau}^{\ell N_2}(\mathbf{z}) \\
 &\leq \epsilon_2/4 + \hat{p}_{s+t,1} + \epsilon_2/4,
 \end{aligned}$$

12 a contradiction. This shows that the probability of \mathbf{z} being good is at least $\epsilon_2/4$. □

```

1: function  $\mathcal{A}_7$ 
2:   for  $N_1$  iterations do
3:     Choose  $t$  uniformly from  $R_q^\vee$ 
4:      $\text{cnt}_0 \leftarrow 0$ ;  $\text{cnt}_1 \leftarrow 0$ ;
5:     for  $N_2$  iterations do
6:       Obtain a list  $\mathcal{T}$  of  $\ell$  samples from  $A_{s,r}^i$ 
7:       Choose  $\mathbf{f}_1, \dots, \mathbf{f}_\ell$  independently from  $D_{\tau}^{\ell}$ 
8:       Compute  $\mathcal{T}'$  and  $\mathcal{T}''$  from  $\mathcal{T}$ ,  $t$  and  $\mathbf{f}_1, \dots, \mathbf{f}_\ell$ 
9:        $\text{cnt}_0 \leftarrow \text{cnt}_0 + \mathcal{D}_1(\mathcal{T}')$ ;  $\text{cnt}_1 \leftarrow \text{cnt}_1 + \mathcal{D}_1(\mathcal{T}'')$ 
10:    end for
11:     $\hat{\mathbf{p}}_0 \leftarrow \text{cnt}_0/N_2$ ;  $\hat{\mathbf{p}}_1 \leftarrow \text{cnt}_1/N_2$ 
12:    if  $|\hat{\mathbf{p}}_0 - \hat{\mathbf{p}}_1| \geq \epsilon_2/4$  then return  $i$ 
13:  end for
14:  return  $i$ 
15: end function.

```

Table 2: Pseudo-code of algorithm \mathcal{A}_7 from distinguisher \mathcal{D}_1 described in Section 4.4.

13 Next we consider the effect of changing the error distribution from $D_{\tau}^{\ell N_2}$ to $D_{\tau'}^{\ell N_2}$. To do this, we introduce
 14 a quantity whose logarithm is the Rényi divergence of order 2. Let k be a positive integer. For two probability

1 density functions²⁰ $P, Q : H^k \rightarrow \mathbb{R}_{\geq 0}$, let

$$\mathbf{R}(P||Q) = \int_{H^k} \frac{P(x)^2}{Q(x)} dx. \quad (29)$$

2 By an abuse of notation, we will also write $\mathbf{R}(D||D')$ to denote $\mathbf{R}(P||Q)$, where D and D' are the distributions
3 corresponding to P and Q respectively. For a measurable subset B of H^k , we have,

$$(\Pr_D[B])^2 = \left(\int_B P(x) dx \right)^2 \quad (30)$$

$$\leq \left(\int_B \frac{P(x)^2}{Q(x)} dx \right) \left(\int_B Q(x) dx \right) \quad (31)$$

$$\leq \left(\int_{H^k} \frac{P(x)^2}{Q(x)} dx \right) \Pr_{D'}[B] \\ = \mathbf{R}(D||D') \Pr_{D'}[B]. \quad (32)$$

4 The derivation of (31) from (30) is made using the Cauchy-Scharwz inequality²¹.

5 **Proposition 11.** *The minimum value of c such that $x^2/\sqrt{2x^2-1}$ is less than $1 + c(x-1)^2$ for $x > 1$ is $c = 2$.*

6 *Proof.* We first show that for $x > 1$, $x^2/\sqrt{2x^2-1} < 1 + 2(x-1)^2$. Let $p(x) = (2x^2-1)(1+2(x-1)^2)^2 - x^4$.
7 Then $x^2/\sqrt{2x^2-1} < 1 + 2(x-1)^2$ if and only if $p(x) > 0$. The polynomial $p(x)$ factors as $p(x) = (8x^3 - 8x^2 +$
8 $3x + 9)(x-1)^3 = (8x^2(x-1) + 3x + 9)(x-1)^3$ which is a sum and product of positive numbers for $x > 1$. Hence,
9 it follows that $p(x) > 0$ for $x > 1$.

10 We next show that if the 2 in $1+2(x-1)^2$ is replaced by $c < 2$, then the inequality $x^2/\sqrt{2x^2-1} < 1+c(x-1)^2$
11 cannot hold when x is close to 1. We set $\epsilon = x - 1$ and $t = 4\epsilon + 2\epsilon^2$ and use the Taylor series $(1+t)^{-1/2} =$
12 $1 - t/2 + 3t^2/8 \pm O(t^3)$. We have

$$\frac{x^2}{\sqrt{2x^2-1}} = (1+\epsilon)^2(1+4\epsilon+2\epsilon^2)^{-1/2} \\ = (1+2\epsilon+\epsilon^2)(1-2\epsilon-\epsilon^2+3(4\epsilon+2\epsilon^2)^2/8 \pm O(\epsilon^3)) \\ = 1+2\epsilon^2 \pm O(\epsilon^3).$$

13 For $c < 2$ and small ϵ , this is greater than $1 + c\epsilon^2$. □

14 **Proposition 12.** *Let $k \geq 1$ be a positive integer and $r_1, r_2 \in \mathbb{R}^+$ be such that $1 < r_2/r_1 < 1 + \sqrt{\ln(nk)/(nk)}/2$.
15 Let D_{r_1} and D_{r_2} be the continuous Gaussian distributions on H having widths r_1 and r_2 respectively. Then*

$$\mathbf{R}(D_{r_1}^k || D_{r_2}^k) \leq \left(1 + \frac{1}{2} \cdot \frac{\ln(nk)}{nk} \right)^{nk}. \quad (33)$$

16 *Proof.* Direct calculation from the definition of the continuous Gaussian distribution D_r on H shows that for
17 $r > 0$ and $x > 1/\sqrt{2}$, $\mathbf{R}(D_r || D_{xr}) = (x^2/\sqrt{2x^2-1})^n$. For $x > 1$, from Proposition 11, we have $(x^2/\sqrt{2x^2-1})^n$
18 is smaller than $(1+2(x-1)^2)^n$. So $\mathbf{R}(D_{r_1}^k || D_{r_2}^k) = (\mathbf{R}(D_{r_1} || D_{r_2}))^k \leq (1 + \ln(nk)/(2nk))^{nk}$. □

²⁰In Claim 5.11 of [19], the density functions are considered to be over \mathbb{R}^n . Here we consider density functions over H^k .

²¹In the Cauchy-Scharwz inequality of the form $(\int_B f(x)g(x)dx)^2 \leq (\int_B f(x)^2dx) (\int_B g(x)^2dx)$, take $f(x) = P(x)/\sqrt{Q(x)}$ and $g(x) = \sqrt{Q(x)}$.

1 Setting $x = \alpha^2/\tau^2$ in the inequality $1 + x < (1 + x/2)^2$ for $x \neq 0$ and using (11), we have

$$1 \leq \frac{r'}{\tau} = \frac{\sqrt{r^2 + \tau^2}}{\tau} \leq \sqrt{1 + \frac{\alpha^2}{\tau^2}} < 1 + \frac{1}{2} \cdot \frac{\alpha^2}{\tau^2} = 1 + \frac{1}{2} \left(\frac{\ln(nN_2\ell)}{nN_2\ell} \right)^{1/2}. \quad (34)$$

2 Applying Proposition 12 with $k = N_2\ell$, $r_1 = \tau$ and $r_2 = r'$, we obtain

$$\mathbb{R}(D_\tau^{\ell N_2} \| D_{r'}^{\ell N_2}) \leq (1 + \ln(nN_2\ell)/(2nN_2\ell))^{nN_2\ell}. \quad (35)$$

3 **Proposition 13.** *The right hand side of (35) is about $(nN_2\ell)^{1/2}$.*

4 *Proof.* The approximation can be seen by setting $x = (2nN_2\ell)/(\ln(nN_2\ell))$ and $m = (\ln(nN_2\ell))/2$ in $(1+1/x)^{mx} \approx$
5 e^m . \square

6 **Remark 22.** *If we follow the analysis in [19], then we obtain $(nN_2\ell)^3$ instead of $(nN_2\ell)^{1/2}$. The net effect of*
7 *this reduction of the exponent on the overall tightness gap is that it is about the square root of the overall tightness*
8 *gap that would be obtained by following the analysis in [19].*

9 **Proposition 14.** *For a good $s+t$, the measure of the set of good \mathbf{z} under $D_{r'}^{\ell N_2}$ is at least about $\epsilon_2^2/(256nN_2\ell)^{1/2}$.*

10 *Proof.* From Section 4.4, we have that for a good $s+t$ the measure of the set of good \mathbf{z} under $D_\tau^{\ell N_2}$ is at least
11 $\epsilon_2/4$. In (32), considering B to be the set of good \mathbf{z} and replacing k by $N_2\ell$, we have

$$\Pr_{D_{r'}^{\ell N_2}}[B] \geq \frac{\left(\Pr_{D_\tau^{\ell N_2}}[B]\right)^2}{\mathbb{R}(D_\tau^{\ell N_2} \| D_{r'}^{\ell N_2})} \geq \frac{\epsilon_2^2}{16\mathbb{R}(D_\tau^{\ell N_2} \| D_{r'}^{\ell N_2})} \gtrsim \frac{\epsilon_2^2}{(256nN_2\ell)^{1/2}}.$$

12

\square

13 **Remark 23.** *In [19], the ratio τ/α is defined to be $((n\ell)/\ln(n\ell))^{1/4}$. If we use this definition of τ/α , and take*
14 *$k = \ell$ in Proposition 12, then instead of (35) we would obtain*

$$\mathbb{R}(D_\tau^{\ell N_2} \| D_{r'}^{\ell N_2}) \leq (\mathbb{R}(D_\tau^\ell \| D_{r'}^\ell))^{N_2} \approx (n\ell)^{N_2/2}. \quad (36)$$

15 *Since $N_2 > n^2$, this would lead to super-exponential running time $> n^{n^2}$.*