Worst-Case to Average-Case Reductions for Module Lattices

Adeline Langlois[1,3] and Damien Stehlé[2,3]

[1] ENS Cachan - Antenne de Bretagne, France
[2] CNRS
[3] Laboratoire LIP (U. Lyon, CNRS, ENS Lyon, INRIA, UCBL),
46 Allée d'Italie, 69364 Lyon Cedex 07, France

Abstract. Most lattice-based cryptographic schemes are built upon the assumed hardness of the Short Integer Solution (SIS) and Learning With Errors (LWE) problems. Their efficiencies can be drastically improved by switching the hardness assumptions to the more compact Ring-SIS and Ring-LWE problems. However, this change of hardness assumptions comes along with a possible security weakening: SIS and LWE are known to be at least as hard as standard (worst-case) problems on euclidean lattices, whereas Ring-SIS and Ring-LWE are only known to be as hard as their restrictions to special classes of ideal lattices, corresponding to ideals of some polynomial rings. In this work, we define the Module-SIS and Module-LWE problems, which bridge SIS with Ring-SIS, and LWE with Ring-LWE, respectively. We prove that these average-case problems are at least as hard as standard lattice problems restricted to module lattices (which themselves generalize arbitrary and ideal lattices). As these new problems enlarge the toolbox of the lattice-based cryptographer, they could prove useful for designing new schemes. Importantly, the worst-case to average-case reductions for the module problems are (qualitatively) sharp, in the sense that there exist converse reductions. This property is not known to hold in the context of Ring-SIS/Ring-LWE: Ideal lattice problems could reveal easy without impacting the hardness of Ring-SIS/Ring-LWE.

1 Introduction

A euclidean lattice is the set of all integer linear combinations of some n linearly independent vectors belonging to a Euclidean space. There are many algorithmic problems related to lattices. In this work, we will consider the *Shortest Independent Vectors problems* (SIVP): The goal is to find n linearly independent vectors $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n$ in a given n-dimensional lattice, that minimize $\max_i \|\boldsymbol{s}_i\|$. A standard relaxation of this optimization problem, parametrized by $\gamma(n) \geq 1$, consists in requesting that $\max_i \|\boldsymbol{s}_i\|$ is within a factor γ of the optimal value. This variant is referred to as SIVP_γ and γ is called the approximation factor. SIVP_γ is known to be NP-hard for any approximation factor $\gamma \leq O(1)$ (see [3]). A standard and well accepted conjecture is to assume that there is no polynomial time algorithm that achieves an approximation factor that is polynomial in n, even using quantum computing [22].

Lattice-based cryptography is a branch of cryptography exploiting the presumed (worst-case) hardness of lattice problems such as SIVP_γ. Its main advantages are its simplicity, efficiency, and apparent security against quantum computers. But perhaps the most appealing aspect is that lattice-based cryptographic protocols often enjoy very strong security proofs based on the hardness of worst-case problems. Typically, an average-case problem (solvers of which correspond to a protocol attacker) is shown to be at least as hard as the arbitrary instances of another problem which is presumed difficult. We refer to [22] for a recent survey on lattice-based cryptography.

Two main problems serve as the foundation of numerous lattice-based cryptographic protocols. The first one, introduced by Ajtai in 1996 [1], is the *Short Integer Solution problem* (SIS): For parameters n, m and q positive integers, the problem is to find a short non zero solution $\boldsymbol{z} \in \mathbb{Z}^m$ to the homogeneous linear system $\boldsymbol{A}\boldsymbol{z} = 0 \bmod q$ for uniformly random $\boldsymbol{A} \in \mathbb{Z}_q^{n \times m}$ (the notation $\mathbb{Z}_q$ denotes the ring of integers modulo q). The second one, introduced by Regev in 2005 [30], is the *Learning With Errors problem* (LWE). The search version of LWE is as follows: For parameters n and q positive integers and χ a probability density function on $\mathbb{T} = \mathbb{R}/\mathbb{Z} \simeq [0, 1)$, the problem is to find $\boldsymbol{s}$, given arbitrarily many independent pairs $(\boldsymbol{a}, \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s} \rangle + e)$

for a vector $\boldsymbol{a} \in \mathbb{Z}_q^n$ chosen uniformly at random, and $e \in \mathbb{T}$ sampled from χ. It is possible to interpret LWE in terms of linear algebra: If m independent samples $(\boldsymbol{a}_i, \frac{1}{q}\langle \boldsymbol{a}_i, \boldsymbol{s}\rangle + e_i)$ are considered, the goal is to find $\boldsymbol{s}$ from $(\boldsymbol{A}, \frac{1}{q}\boldsymbol{A}\boldsymbol{s} + \boldsymbol{e})$, where the rows of $\boldsymbol{A}$ correspond to the $\boldsymbol{a}_i$'s and $\boldsymbol{e} = (e_1, \ldots, e_m)^T$. The decision counterpart of LWE consists in distinguishing between arbitrarily many independent pairs $(\boldsymbol{a}, \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s}\rangle + e)$ sampled as in the search version and the same number of uniformly random and independent pairs.

Ajtai [1] proposed the first worst-case to average-case reduction for a lattice problem, by providing a reduction from SIVP_γ to SIS. Later, Regev [30,31] showed the hardness of the LWE problem by describing a (quantum) reduction from SIVP_γ to LWE. Cryptographic protocols relying on SIS or LWE therefore enjoy the property of being provably as secure as a worst-case problem which is strongly suspected of being extremely hard. However, on the other hand, the cryptographic applications of SIS and LWE are inherently inefficient due to the size of the associated key (or public data), which typically consists of the matrix $\boldsymbol{A}$.

To circumvent this inherent inefficiency, Micciancio [15,17] — inspired from the efficient NTRU encryption scheme [11] that can be interpreted in terms of lattices — initiated an approach that consists in changing the SIS and LWE problems to variants involving structured matrices. In these variants, the random matrix $\boldsymbol{A}$ is replaced by one with a specific block-Toeplitz structure, thus allowing for more compact keys and more efficient algorithms. The problem considered by Micciancio in [17] was later replaced by a more powerful variant [13,26], now commonly referred to as Short Integer Solution problem over Rings, or R-SIS (it was initially called Ideal-SIS). A similar adaptation for LWE, called R-LWE, was introduced by Lyubashevsky et al [14] (see also [33]). Similarly to SIS and LWE, these problems admit reductions from worst-case lattice problems [13,26,14], but, however, the corresponding worst-case problem is now SIVP_γ restricted to ideal lattices (which correspond to ideals of the ring of integers of a number field corresponding to the specific matrix structure). The latter problem is called Id-SIVP.

Main results. In this paper, we bridge the reductions from SIVP to SIS and Id-SIVP to R-SIS on the first hand, and from SIVP to LWE and Id-SIVP to R-LWE on the second hand. We consider two problems M-SIS and M-LWE, where the letter M stands for module. A module is an algebraic structure generalizing rings and vector spaces, whereas module lattices (corresponding to finitely generated modules over the ring of integers of a number field) generalize both arbitrary lattices and ideal lattices. Note that M-LWE has recently been introduced (although not studied) in [6], where it is called Generalized-LWE. We describe two new worst-case to average-case reductions: A reduction from Mod-SIVP (i.e., SIVP restricted to module lattices) to M-SIS, and a (quantum) reduction from Mod-SIVP to M-LWE in both its search and decision versions.

The Mod-SIVP to M-SIS and Mod-SIVP to M-LWE reductions are smooth generalizations of the existing reductions: By setting the module dimension and the field degree appropriately, we recover the former reductions. When doing so, the conditions on the approximation factor γ and the modulus q required for the results to hold match with the conditions of the existing reductions, up to a factor that is logarithmic in the lattice dimension. These parameters quantify the quality of the reductions: The hardness of the SIVP problem is given by the approximation factor γ, whereas the bit-size of the average-case instances is proportional to $\log q$.

To achieve these results, we carefully combine and adapt the existing reductions and their proofs of correctness ([9] and [13] for M-SIS, and [31] and [14] for M-LWE). At a high level, the module structure can be seen as a "tensor" between the lattice and ideal algebraic structures, leading to reductions and proof that can heuristically be seen as "tensors" of the former reductions and proofs.

A larger toolbox for the lattice cryptographer. The hardness results for M-SIS and M-LWE possibly enlarge the tool box for devising lattice-based cryptosystems. Let us consider small examples. The following is an instance of M-SIS for which we can prove hardness for specific values of the parameters n, q and β. Given $a_{i,j}$'s sampled uniformly and independently from the uniform distribution over $\mathbb{Z}_q[x]/(x^n+1)$, the goal is to find z_i's in $\mathbb{Z}[x]/(x^n+1)$ not all zero, with coefficients smaller than a prescribed bound β and such that:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{bmatrix} \cdot \begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix} = 0 \mod q.$$

Similarly, our results on M-LWE imply that for specific values of n, q and for a specific error distribution ψ taking small values in $\mathbb{Z}[x]/(x^n+1)$ (or, actually, a specific distribution over such distributions), the following pair is computationally indistinguishable) from uniform over its range:

$$\left(\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix}, \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \\ a_{31} & a_{32} \end{bmatrix} \cdot \begin{bmatrix} s_1 \\ s_2 \end{bmatrix} + \begin{bmatrix} e_1 \\ e_2 \\ e_3 \end{bmatrix} \bmod q \right),$$

where the a_{ij}'s and s_i's are sampled uniformly in $\mathbb{Z}_q[x]/(x^n+1)$, and the e_i's are sampled from ψ.

Note that the existing results on R-LWE and R-SIS already imply that these problems are as hard as some SIVP instances: For example, one can embed an R-SIS instance into the first row of an M-SIS instance, and generate the other row(s) independently. However, with this approach, the hardness of the corresponding worst-case instances is related to n-dimensional instances of SIVP. With our new approach, we can show that the M-SIS instance above is as hard as solving SIVP for a $(2n)$-dimensional lattice (or, more generally, a (dn)-dimensional lattice, if the number of rows of the M-SIS matrix is d). If SIVP is exponentially hard to solve (with respect to the lattice dimension), the module approach provides a complexity lower bound for solving this simple M-SIS instance that is the square (resp. dth power) of the lower bound provided by relying on R-SIS.

Hedging against a possible non-hardness of Id-SIVP. The SIVP to SIS and SIVP to LWE reductions are qualitatively sharp in the sense that they allow for converse reductions: Both LWE and SIS can be solved using an SIVP solver. Such a result is not known to hold for Id-SIVP and R-SIS/R-LWE, possibly hinting at a weakness of Id-SIVP. We show (in Section 5) that the Mod-SIVP to M-SIS/M-LWE do not suffer from this drawback. Furthermore, Mod-SIVP can trivially be shown to be no easier than Id-SIVP (as any Id-SIVP instance can be encoded into a Mod-SIVP instance of higher dimension). As a consequence, our results lead to cryptographic primitives whose efficiencies are within a constant factor of those based on R-SIS/R-LWE, but for which the worst-case to average-case reduction is sharp.

Id-SIVP has been much less studied than SIVP, and attacks on SIVP working only in the case of ideal lattices cannot be fully ruled out. Such attacks could, for example, exploit the multiplicative structure of the ideals. Such weaknesses due to the multiplicative structure actually exist in the case for the decisional counterparts of SIVP_γ: It becomes easy for ideal lattices (assuming γ is not too small), because good and efficiently computable estimates are known for the successive minima of any given ideal lattice (see, e.g., [27, Se. 6]). It is an important open problem to assess the hardness of Id-SIVP.

Related works. Most SIVP to SIS reductions (including ours) consider the euclidean norm. Peikert [24] described an SIVP to SIS reduction that handles all ℓ_p norms. Independently, many variants of LWE have been shown as hard as Regev's original LWE: These variants may consist in sampling the secret vector $\boldsymbol{s}$ from the same distribution as the errors [2], in sampling the error vectors from other distributions [25,10] and in relaxing the conditions on the factorisation of the modulus [19, Se. 3] (see also the references therein). In [25], Peikert partially dequantumized Regev's proof of hardness of LWE [31], by proposing a reduction from (several variants of) the decisional GapSVP problem to LWE. GapSVP_γ consists in estimating the minimum of the input lattice to within a multiplicative factor γ. Peikert's classical reduction is restricted to large LWE moduli q (that are additionally required to be products of many small primes in the case of the decisional variant of LWE), unless one considers a variant of GapSVP that is somewhat unusual. Peikert's dequantumization carries over to the module case, by giving a reduction from GapSVP restricted to module lattices to M-LWE (using Lemma 17 from Section 4). Note that it also carries over to ideal/R-LWE setting but is meaningless in this situation as GapSVP is easy for ideal lattices and the involved approximation factors γ (as a good approximation to the minimum known). Other cryptographically useful variants of SIS and LWE proven as secure as SIVP include k-SIS [4], ISIS [9] and subspace-LWE [12,28].

Some computational aspects of module lattices have been investigated in [5,8] (see also [7, Ch. 1]). These results show that the additional algebraic structure may be exploited to obtain compact representations of modules (namely, pseudo-bases) similar to lattice bases in Hermite Normal Form and LLL-reduced lattice bases. None hints that SIVP would be any easier when restricted to module lattices.

Road-map. We first give reminders on euclidean lattices, elementary algebraic number theory and Gaussian measures. In Section 3 we give a reduction from Mod-SIVP to M-SIS. Then, in Section 4, we describe a (quantum) reduction from Mod-SIVP to both the computational and the decisional variants of M-LWE. Finally, we give converse reductions in Section 5, i.e., reductions from both M-SIS and M-LWE to Mod-SIVP.

2 Preliminaries

Notation. Vectors will be denoted in bold, and if $\boldsymbol{x}$ is a vector, then its ith coordinate will be denoted by x_i and its euclidean norm will be denoted by $\|\boldsymbol{x}\|$. For a tuple of vectors $\boldsymbol{X} = (\boldsymbol{x}_i)_i$, we let $\|\boldsymbol{X}\| = \max_i \|\boldsymbol{x}_i\|$. The vector $\boldsymbol{e}_i$ denotes the vector with 1 in its ith coordinate and 0 in all its other coordinates. Let $\boldsymbol{B} \in \mathbb{R}^n$ be a basis. We let $\tilde{\boldsymbol{B}}$ be the Gram-Schmidt orthogonalisation of $\boldsymbol{B}$.

We use standard Landau notations. Furthermore, we say that a function $f(n)$ is $\mathrm{poly}(n)$ if it is bounded by a polynomial in n. The notation $\omega(f(n))$ refers to the set of functions (or an arbitrary function in that set) growing faster than $c \cdot f(n)$ for any constant $c > 0$. A function $\varepsilon(n)$ is said negligible if it decreases faster than the inverse of any polynomial function, i.e., if it is $n^{-\omega(1)}$. Finally, a function is exponentially small in n if it is at most $2^{-\Omega(n)}$.

The statistical distance between two distributions X and Y on a countable set D is defined as follows: $\Delta(X, Y) = \frac{1}{2} \sum_{d \in D} |X(d) - Y(d)|$. We say that two sequences $(X_n)_n, (Y_n)_n$ of distributions indexed by a variable n are negligibly close if $\Delta(X_n, Y_n)$ is negligible in n. Finally, we let $U(S)$ denote the uniform distribution over the set S.

Remark on the reductions. The worst-case lattice problem SIVP_γ is suspected to be exponentially hard to solve with respect to the lattice dimension, even using quantum computations. This is one of its most attractive features for using it as a hardness assumption for devising cryptographic primitives. For this reason, we consider two types of worst-case to average-case reductions from SIVP (and its ideal and module variants): Reductions that assume we have a polynomial time algorithm for solving the considered average-case problem with non-negligible probability, and (somewhat more unusually), reductions that assume we have a subexponential-time algorithm for solving the considered average-case problem with non-exponentially small probability.

2.1 Some algebraic number theory

We briefly recall a few facts on elementary algebraic number theory. We refer the reader to [23] for a thorough introduction.

Number fields. Every complex root of a polynomial $g(X) \in \mathbb{Q}[X]$ is an *algebraic number*. The *minimal polynomial* of an algebraic number ξ is the unique irreducible monic polynomial f of minimal degree such that ξ is one of its roots. An *algebraic integer* is an algebraic number whose minimal polynomial is in $\mathbb{Z}[X]$. Let ξ be an algebraic number, the *number field* $K = \mathbb{Q}(\xi)$ is a finite extension of the rational number field $\mathbb{Q}$. It is also an n-dimensional vector space over $\mathbb{Q}$ with basis $\{1, \xi, \ldots, \xi^{n-1}\}$, where n is the degree of f. We call n the degree of K. An element $x \in K$ can be represented by the coefficients of the associated rational polynomial modulo f: $\sigma_P(x) = (x_i)_i$ is such that $x = \sum_{i=0}^{n-1} x_i \cdot \xi^i$. The product between two elements corresponds to the product between the two associated polynomials modulo f. Let R be the set of the algebraic integers belonging to K. This is a ring, called the ring of integers (or maximal order) of K. If ξ is an algebraic integer, then $\mathbb{Z}[\xi] = \sum_{i=1}^{n} \mathbb{Z} \cdot \xi^i \subseteq R$, but in general this inclusion can be strict.

Complex embeddings. Let $K = \mathbb{Q}(\xi)$ be a degree n number field, and let f be the minimal polynomial of ξ. The canonical embeddings are the n homomorphisms $\sigma_i : K \to \mathbb{C}$ that fix every element of $\mathbb{Q}$. We let $\{\sigma_i\}_{i \in [s_1]}$ denote the real embeddings (corresponding to the real roots of f), and $\{\sigma_i\}_{s_1 < i \leq s_1 + 2s_2}$ denote the complex embeddings (corresponding to the complex roots). We have $s_1 + 2s_2 = n$ and we can reorder the complex embeddings so that $\sigma_{s_1+s_2+i} = \overline{\sigma_{s_1+i}}$ for all $i \in [s_2]$. We call *canonical embedding* the ring

homomorphism $\sigma_C : K \to \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ defined as: $\sigma_C(y) = (\sigma_1(y), \ldots, \sigma_n(y))$. An element of K is fully determined by its canonical embedding. For any $x, y \in K$, we have that $\sigma_C(x \cdot y)$ is the component wise product of $\sigma_C(x)$ and $\sigma_C(y)$.

The *trace* $\mathrm{Tr} : K \to \mathbb{Q}$ and the *norm* $\mathrm{N} : K \to \mathbb{Q}$ are defined as follows: $\mathrm{Tr}(x) = \sum_{i\in[n]} \sigma_i(x)$ and $\mathrm{N}(x) = \prod_{i\in[n]} \sigma_i(x)$. For any $x, y \in K$ we have $\mathrm{Tr}(x \cdot y) = \sum_{i\in[n]} \sigma_i(x) \cdot \sigma_i(y) = \langle \sigma_C(x), \overline{\sigma_C(y)} \rangle$ where $\langle \cdot, \cdot \rangle$ is the canonical Hermitian product on $\mathbb{C}^n$.

Space H. As in [14], we will use the following subspace of $\mathbb{C}^n$:

$$H = \{(x_1, \ldots, x_n) \in \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2} : \forall j \in [s_2], x_{s_1+s_2+j} = \overline{x_{s_1+j}}\}.$$

The space H is isomorphic to $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$. Let $\boldsymbol{h}_j = \boldsymbol{e}_j$ for $j \in [s_1]$ and $\boldsymbol{h}_j = \frac{1}{\sqrt{2}}(\boldsymbol{e}_j + \boldsymbol{e}_{j+s_2})$ and $\boldsymbol{h}_{j+s_2} = \frac{i}{\sqrt{2}}(\boldsymbol{e}_j - \boldsymbol{e}_{j+s_2})$ for $s_1 < j \leq s_1 + s_2$. The $\boldsymbol{h}_j$'s form a basis of H as a real vector space. An element $x \in K$ can be represented according to the basis $(\boldsymbol{h}_j)_j$: We define $\sigma_H(x)$ by $\sigma_H(x) = (x_j)_j \in \mathbb{R}^n$ such that $\sigma_C(x) = \sum_{j=1}^n x_j \cdot \boldsymbol{h}_j$.

Ideals. An (integral) *ideal* I of R is an additive subgroup of R that is closed under multiplication by every element of R. The smallest ideal of R containing the set S is denoted by (S). The quotient R/I is the set of the equivalence classes $g + I$ of R modulo I. For any non-zero ideal, the *norm* $N(I)$ of the ideal is the number of elements of the quotient ring R/I. We have $N((x)) = \mathrm{N}(x)$, for all $x \in K$.

Let I and J be ideals of R. We define the *product* of two ideals by $IJ = \{\sum_i \alpha_i\beta_i : \alpha_i \in I, \beta_i \in J\}$ and their *sum* by $I + J = \{\alpha + \beta : \alpha \in I, \beta \in J\}$. The ideals I and J are said *coprime* if $I + J = R$. As we have $(I + J)(I \cap J) = IJ$, if I and J are coprime, then $I \cap J = IJ$. An ideal $I \subsetneq R$ is *prime* if for any $ab \in I$ then $a \in I$ or $b \in I$. In R, an ideal I is prime if and only if it is *maximal*, which implies that the quotient ring R/I is the finite field of order $N(I)$. Finally, every ideal of R can be represented as a unique product of prime ideals.

A *fractional ideal* $I \subseteq K$ is a set such that $dI \subseteq R$ is an (integral) ideal for a non-zero $d \in R$. The *inverse* of a fractional I is defined by $I^{-1} = \{\alpha \in K : \alpha I \subseteq R\}$ and is itself a fractional ideal. We have $II^{-1} = R$. The *dual* of an ideal is defined as $I^* = \{x \in K : \mathrm{Tr}(xI) \subseteq \mathbb{Z}\}$. We have $I^* = I^{-1} \cdot R^*$.

Isomorphisms of quotient rings. Lyubashevsky et al [14, Se. 2.3.7] made explicit an isomorphism between I/qI and R/qR for an arbitrary positive integer q, which we recall now. Let R_q and R_q^* respectively denote R/qR and R^*/qR^*.

Let $t \in I$ be such that $(t)+qI = I$ (such a t exists and can be found efficiently given I and the prime ideal factorization of (q), see [14, Le. 2.8]). The function $\theta_I : K \to K$ defined as $\theta_I(x) = t \cdot x$ induces an isomorphism from R_q to I/qI. Moreover, this isomorphism may be efficiently inverted using $\theta_I^{-1} : I/qI \to R_q$ defined by $\theta_I^{-1}(y) = t^{-1} \cdot y' \bmod qR$ where $y' = y \bmod qI$ and $y' \in (t)$. The function θ_I also induces an isomorphism from I^*/qI^* to R_q^* that may be efficiently inverted using $\theta_I^{-1} : R_q^* \to I^*/qI^*$ with $\theta_I^{-1}(y) = t^{-1} \cdot y' \bmod qR$ where $y' = y \bmod qI^*$ and $y' \in (t)$.

Modules. A subset $M \subseteq K^d$ is a module if it is closed under addition and multiplication by elements of R. It is a finitely generated module if there exists a finite family $(\boldsymbol{b}_i)$ such that $M = \sum_i R \cdot \boldsymbol{b}_i$. In general, if the ring R is arbitrary, an R-module may not have a basis. But here K is a number field, so R is a Dedekind domain, and we have the existence of so-called pseudo-bases (see, e.g., [7, Ch. 1]): For every module M, there exist I_i ideals of R and $(\boldsymbol{b}_i)_i$ linearly independent vectors of K^d such that $M = \sum_{i=1}^d I_i \cdot \boldsymbol{b}_i$. We say that $[(I_i)_i, (\boldsymbol{b}_i)_i]$ is a *pseudo-basis* of M. The representation of the elements of M with respect to a pseudo-base is unique. Two pseudo-bases can generate the same module and then, they have the same cardinal. The latter is called *rank* of the module.

We define the dual of a module by $M^* = \{\boldsymbol{y} \in K^d, \forall \boldsymbol{x} \in M : \mathrm{Tr}(\langle \boldsymbol{x}, \boldsymbol{y} \rangle) \in \mathbb{Z}\}$, where $\langle \cdot, \cdot \rangle$ is the canonical inner product on K^d. We have the following property:

Lemma 1. *If $M = \sum_{i=1}^d I_i \cdot \boldsymbol{b}_i$, then $M^* = \sum_{i=1}^d I_i^* \cdot \boldsymbol{b}_i^*$, where the $\boldsymbol{b}_i^*$'s are defined by $\forall i, j$, $\langle \boldsymbol{b}_i, \boldsymbol{b}_j^* \rangle = 1$ if $i = j$ and $\langle \boldsymbol{b}_i, \boldsymbol{b}_j^* \rangle = 0$ otherwise.*

Proof. We first show that $\sum_{i=1}^{d} I_i^* \cdot \boldsymbol{b}_i^* \subseteq M^*$. Let $\boldsymbol{y} \in \sum_{i=1}^{d} I_i^* \cdot \boldsymbol{b}_i^*$. Then for each i there exists $y_i \in I_i^*$ such that $\boldsymbol{y} = \sum_{i=1}^{d} y_i \cdot \boldsymbol{b}_i^*$. Let $\boldsymbol{x} = \sum_{i=1}^{d} x_i \cdot \boldsymbol{b}_i \in M$. Then by linearity, we have $\mathrm{Tr}(\langle \boldsymbol{x}, \boldsymbol{y} \rangle) = \sum_{i=1}^{d} \mathrm{Tr}(x_i y_i)$. For all i, we have $x_i \in I_i$ and $y_i \in I_i^*$, and thus $\mathrm{Tr}(x_i y_i) \in \mathbb{Z}$. Therefore, we have $\mathrm{Tr}(\langle \boldsymbol{x}, \boldsymbol{y} \rangle) \in \mathbb{Z}$ and $\boldsymbol{y} \in M^*$.

We now show that $M^* \subseteq \sum_{i=1}^{d} I_i^* \cdot \boldsymbol{b}_i^*$. Let $\boldsymbol{y} \in M^* \subset K^d$. We can write $\boldsymbol{y} = \sum_{i=1}^{d} y_i \cdot \boldsymbol{b}_i^*$, for some y_i's in K. It suffices to show that $y_i \in I_i^*$. By linearity, we have $\mathrm{Tr}(\langle \boldsymbol{y}, x_i \boldsymbol{b}_i \rangle) = \mathrm{Tr}(x_i y_i)$. And we obtain $\mathrm{Tr}(\langle \boldsymbol{y}, x_i \boldsymbol{b}_i \rangle) \in \mathbb{Z}$. This implies that $y_i \in I_i^*$. □

We generalize the isomorphism θ_I defined above to modules. Let $M = \sum_{i=1}^{d} I_i \cdot \boldsymbol{b}_i$, $f : I_1/qI_1 \times \ldots \times I_d/qI_d \to M/qM$ be such that $f(x_1, \ldots, x_n) = \sum_{i=1}^{d} x_i \cdot \boldsymbol{b}_i$ and $g : M/qM \to I_1/qI_1 \times \ldots \times I_d/qI_d$ be such that $g(\sum_{i=1}^{d} x_i \cdot \boldsymbol{b}_i) = (x_1, \ldots, x_n)$. The functions f and g are ring isomorphisms and $g = f^{-1}$. Let $\theta_{I_1}, \ldots, \theta_{I_d}$ be as described above. We define the functions Θ and Θ^{-1} as follows: $\Theta = f \circ (\theta_{I_1} \times \ldots \times \theta_{I_d})$ and $\Theta^{-1} = (\theta_{I_1}^{-1} \times \ldots \times \theta_{I_d}^{-1}) \circ g$. The function Θ induces an isomorphism from $(R_q)^d$ to M/qM with inverse Θ^{-1}.

Cyclotomic fields. A cyclotomic field is a field $K = \mathbb{Q}(\xi)$ where ξ is a primitive root of unity. In this work, we will only consider fields of the form $\mathbb{Q}(\xi)$ for ξ a primitive 2^k-th root of unity.[4] In this setup, there is an isomorphism between K and $\mathbb{Q}[x]/\langle x^n + 1 \rangle$ with $n = 2^{k-1}$. We have that $s_1 = 0$ and $s_2 = n/2$ because none of the roots of $f(x) = x^n + 1$ is real.

This choice of field K leads to several simplifying properties. First, we have $R = \mathbb{Z}[\xi]$. The isomorphism $K \simeq \mathbb{Q}[x]/\langle x^n + 1 \rangle$ induces an isomorphism between R and $\mathbb{Z}[x]/\langle x^n + 1 \rangle$. Second, the representations of the elements of K introduced above (σ_P, σ_C and σ_H) are equivalent up to a similarity. In particular, the ratios $\frac{\sigma_P(x)}{\sigma_C(x)}$ and $\frac{\sigma_P(x)}{\sigma_H(x)}$ are constant with respect to $x \in K \setminus \{0\}$. Another simplifying property is that $R^* = \frac{1}{n} R$.

If q is prime, the prime ideal factorization of $(q) \subseteq R$ can be computed efficiently. In particular, if $q = 1 \bmod 2n$, then $(q) = \prod_{i=1}^{n} \mathfrak{q}_i$ where each $\mathfrak{q}_i$ if a prime ideal with norm $N(\mathfrak{q}_i) = q$. Finally, the field K has $\varphi(2n) = n$ automorphisms $\tau_k : K \to K$ defined by $\tau_k(\xi) = \xi^k$ (for $1 \leq k \leq n$). As noted in [14, Le. 2.10], the automorphism group of the $\{\tau_k\}$ acts transitively on the set $\{\mathfrak{q}_i\}_{i \leq n}$.

2.2 Lattices

We refer to [18,29] for introductions to lattices and their computational aspects. A *euclidean lattice* $\Lambda \subseteq \mathbb{R}^n$ is the set of all integer linear combinations $\sum_{i=1}^{p} \beta_i \boldsymbol{b}_i$ of some linearly independent vectors $(\boldsymbol{b}_i)_{1 \leq i \leq p} \in \mathbb{R}^n$. We write $\mathcal{L}(\boldsymbol{B})$ for the lattice spanned by the basis $\boldsymbol{B} = (\boldsymbol{b}_i)_{i \leq p}$. We call p the dimension of the lattice. In this work, we will restrict ourselves to full-rank lattices, i.e., with $p = n$.

The *minimum* $\lambda_1(\Lambda)$ of a lattice Λ is the norm of any of its shortest non-zero vectors. More generally, the *ith successive minimum* $\lambda_i(\Lambda)$ is the smallest radius r such that Λ contains i linearly independent vectors of norm at most r. The *dual lattice* of $\Lambda \subseteq \mathbb{R}^n$ is $\Lambda^* = \{\boldsymbol{x} \in \mathbb{R}^n : \forall \boldsymbol{y} \in \Lambda, \langle \boldsymbol{x}, \boldsymbol{y} \rangle \in \mathbb{Z}\}$. If $\Lambda = \mathcal{L}(\boldsymbol{B})$ then $\Lambda^* = \mathcal{L}(\boldsymbol{B}^*)$ with $\boldsymbol{B}^* = \boldsymbol{B}^{-T}$.

We consider the following generalization of SIVP. Let ϕ denote an arbitrary real-valued function of a lattice (e.g., taking $\phi = \lambda_n$ allows one to recover SIVP_γ). Let $\gamma \geq 1$ be a function of the dimension n. The *Generalized Independent Vectors Problem* $\mathrm{GIVP}_\gamma^\phi$ is as follows: Given a lattice basis $\boldsymbol{B}$, find $n = \dim(\mathcal{L}(\mathbf{B}))$ linearly independent vectors $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_n \in \mathcal{L}(\mathbf{B})$ such that $\max_i \|\boldsymbol{s}_i\| \leq \gamma \cdot \phi(\mathcal{L}(\mathbf{B}))$.

For $\phi = \lambda_n$, this problem is NP-hard for any approximation factor $\gamma \leq O(1)$ (see [3]). The best known algorithms (even quantum) for an exact solution and an approximation to within any polynomial factor γ all have exponential complexities [22]. This motivates the following conjecture: There is no polynomial time (quantum) algorithm that approximates lattice problems to within a polynomial factor. The following stronger conjecture also seems to hold: There is no sub-exponential time (quantum) algorithm that approximates lattice problems to within a polynomial factor.

[4] Our results can be generalized to all cyclotomic fields, but we restrict ourselves to these ones for the sake of simplicity.

Ideal and module lattices. Let ϕ be an embedding from K to $\mathbb{R}^n$ and I an ideal of R. Then $\phi(I)$ is a lattice. We call it ideal lattice with respect to K and ϕ. We will only consider $\phi = \sigma_P$ of $\phi = \sigma_H$. We will use the first representation in context of M-SIS and the second one in context of M-LWE. In the case of $\phi = \sigma_P$, notice that $\|\sigma_P(\xi x)\| = \|\sigma_P(x)\|$ for any $x \in K$. As a consequence, we have $\lambda_1(\Lambda) = \lambda_n(\Lambda)$ for any n-dimensional ideal lattice Λ. We let Id-GIVP denote the restriction of GIVP to ideal lattices.

We define module lattices similarly. Let ϕ be an embedding from K^d to $\mathbb{R}^{nd}$ and $M \subseteq K^d$ a module of R, then $\phi(M)$ is a module lattice. We will consider in particular $\phi = (\sigma_P, \ldots, \sigma_P)$ and $\phi = (\sigma_H, \ldots, \sigma_H)$. To ease the presentation we call them σ_P and σ_H respectively. Similarly to ideal lattices, we let Mod-GIVP denote the restriction of GIVP to ideal lattices. Note that if M is a rank d module and if K has degree n, then the corresponding module lattice has dimension nd. When a module is given as input of a problem, we consider that we give a lattice basis of the corresponding module lattice. Note that it is equivalent to give a basis of the module lattice and a pseudo-basis of the module because from the first representation, the second representation is computable in polynomial time [5,7]. All asymptotic statements involving modules (including hardness results) will be given for nd growing to infinity.

2.3 Gaussian measures

For a vector $\boldsymbol{c} \in \mathbb{R}^n$ and a real $s > 0$, the Gaussian function is defined by $\rho_{s,\boldsymbol{c}}(\boldsymbol{x}) = e^{-\pi\|\frac{\boldsymbol{x}-\boldsymbol{c}}{s}\|^2}$, for all $\boldsymbol{x} \in \mathbb{R}^n$. This function is extended to any countable set $A \subseteq \mathbb{R}^n$ in the usual way: $\rho_{s,\boldsymbol{c}}(A) = \sum_{\boldsymbol{x}\in A} \rho_{s,\boldsymbol{c}}(\boldsymbol{x})$. By normalizing the Gaussian function, we obtain the continuous Gaussian probability distribution: $\nu_s(\boldsymbol{x}) = \rho_s(\boldsymbol{x})/s^n$. For $\boldsymbol{r} = (r_1, \ldots, r_n) \in (\mathbb{R}^+)^n$, a sample from $\nu_{\boldsymbol{r}}$ over $K_{\mathbb{R}}$ is given by $(\nu_{r_i})_i$. We define $\Psi_{\leq\alpha}$ for $\alpha > 0$, as the set of Gaussian distributions $\nu_{\boldsymbol{r}}$ with $r_i < \alpha$, for all i. For all $\boldsymbol{c} \in \mathbb{R}^n$, $s > 0$ and lattice Λ, the discrete Gaussian probability distribution is defined by:

$$\forall \boldsymbol{x} \in \Lambda,\ D_{\Lambda,s,\boldsymbol{c}}(\boldsymbol{x}) = \frac{\rho_{s,\boldsymbol{c}}(\boldsymbol{x})}{\rho_{s,\boldsymbol{c}}(\Lambda)}.$$

The following theorem ensures that for s large enough, it is possible to efficiently sample according to a discrete Gaussian distribution.

Theorem 1 ([9, Th. 4.1]). *There is a probabilistic polynomial time algorithm that, given a basis $\boldsymbol{B}$ of an n-dimensional lattice $\Lambda = \mathcal{L}(\boldsymbol{B})$, a standard deviation $s \geq \|\tilde{\boldsymbol{B}}\| \cdot \omega(\sqrt{\log n})$ (resp. $s \geq \|\tilde{\boldsymbol{B}}\| \cdot \Omega(\sqrt{n})$), and a center $\boldsymbol{c} \in \mathbb{R}^n$, outputs a sample that is negligibly (resp. exponentially) close to $D_{\Lambda,s,\boldsymbol{c}}$.*

The *smoothing parameter* of a lattice was introduced by [21]. For an n-dimensional lattice Λ and a positive real $\varepsilon > 0$, the smoothing parameter $\eta_\varepsilon(\Lambda)$ is the smallest s such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$. This parameter gives a threshold above which many properties for continuous Gaussians also carry over to discrete Gaussians. We recall a few standard properties on discrete Gaussians that we will need in our reductions.

Lemma 2 ([21, Le. 3.3]). *Let Λ be an n-dimensional lattice and $\varepsilon > 0$. Then $\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}} \cdot \lambda_n(\Lambda)$.*

The latter result implies a (trivial) reduction from SIVP_γ to $\mathrm{GIVP}^{\eta_\varepsilon}_{\gamma'}$, with $\gamma' = \gamma/\sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}}$. We will describe worst-case to average-case reductions involving $\mathrm{GIVP}^{\eta_\varepsilon}$ instead of SIVP. In our reductions, we will consider two choices for ε: For the polynomial time reductions, we will use $\varepsilon = n^{-\omega(1)}$, and for the sub-exponential time reductions we will take $\varepsilon = 2^{-\Omega(n)}$.

Lemma 3 ([24, Le. 3.5]). *Let Λ be an n-dimensional lattice and $\varepsilon > 0$. Then $\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}}/\lambda_1^\infty(\Lambda^*)$, where λ_1^∞ refers to the minimum with respect to the infinity norm.*

Lemma 4 ([9, Cor. 2.8]). *Let $\Lambda' \subseteq \Lambda$ be n-dimensional lattices. Then for any $\varepsilon \in (0,1)$, any $s \geq \eta_\varepsilon(\Lambda')$, and any $\boldsymbol{c} \in \mathbb{R}^n$, the distribution $(D_{\Lambda,s,\boldsymbol{c}} \bmod \Lambda')$ is within statistical distance at most 2ε of the uniform distribution over Λ/Λ'.*

Lemma 5 ([20, Le. 4.4]). *Let Λ be an n-dimensional lattice, $s > 2\eta_\varepsilon(\Lambda)$ for $\varepsilon \leq 1/100$, and $\boldsymbol{c} \in \mathbb{R}^n$. Then for any $(n-1)$-dimensional hyperplane $\mathcal{H}$, the probability that $x \notin \mathcal{H}$ where x is chosen from $D_{\Lambda,s,c}$ is $\geq 1/100$.*

2.4 Linear combinations of Gaussians

The sum of two continuous Gaussians with parameters s and r is a continuous Gaussian with parameter $\sqrt{s^2+r^2}$. We have the following lemma for the sum of a continuous Gaussian and a discrete one.

Lemma 6 ([31, Claim 3.9]). *Let Λ be a lattice, $\boldsymbol{u} \in \mathbb{R}^n$, $r, s > 0$ and $t = \sqrt{r^2+s^2}$. Assume that $rs/t \geq \eta_\varepsilon(\Lambda)$ for some $\varepsilon < 1/2$. Consider the continuous distribution Y on $\mathbb{R}^n$ obtained by sampling from $D_{\Lambda+\boldsymbol{u},r}$ and then adding a vector taken from ν_s. Then the statistical distance between Y and ν_t is at most 4ε.*

We will also use the following result:

Lemma 7 (Adapted from [24, Cor. 5.3]). *For any lattice $\Lambda \subseteq \mathbb{R}^n$, $\boldsymbol{c} \in \mathbb{R}^n$, $\varepsilon \in (0,1)$, $t \geq \sqrt{2\pi}$, unit vector $\boldsymbol{u} \in \mathbb{R}^n$ and $s \geq \eta_\varepsilon(\Lambda)$, we have:*

$$\Pr_{\boldsymbol{b} \hookleftarrow D_{\Lambda,s,\boldsymbol{c}}}[|\langle \boldsymbol{b}-\boldsymbol{c}, \boldsymbol{u}\rangle| \geq st] \leq \frac{1+\varepsilon}{1-\varepsilon} t\sqrt{2\pi e} \cdot e^{-\pi t^2}.$$

We generalize [24, Cor. 5.3] and [32, Le. 2.9] to the case of module lattices.

Lemma 8. [5] *Let $\varepsilon \in (0, \frac{1}{2m+1})$ and $(z_1, \ldots, z_m) \in R^m$. Let $M \subseteq K^d$ be a rank d module on R, $s \geq \eta_\varepsilon(M)$ and $(\boldsymbol{c}_1, \ldots, \boldsymbol{c}_m) \in (R^d)^m$. If the $\boldsymbol{y}_i$'s are sampled from $D_{M,s,\boldsymbol{c}_i}$, then for all $t \geq 0$:*

$$\Pr\left[\left\|\sum_{i=1}^m (\boldsymbol{y}_i - \boldsymbol{c}_i)\cdot z_i\right\|_\infty \geq st\|\boldsymbol{z}\|\right] \leq \frac{1+\varepsilon}{1-\varepsilon} tnd\sqrt{2\pi e} \cdot e^{-\pi t^2}.$$

In particular, for $t = \omega(\sqrt{\log nd})$ (resp. $t = \Omega(\sqrt{nd})$) the above probability is negligible (resp. exponentially small) with respect to nd.

Proof. This proof builds upon that of [24, Cor. 5.3].

The principle is to interpret the m Gaussian samples from the nd-dimensional lattice M as one Gaussian sample from the ndm-dimensional lattice L and then apply Lemma 7, where $L = M \times \cdots \times M$ (i.e., the Cartesian product of m copies of M). We also define $\boldsymbol{c}' = (\boldsymbol{c}_1, \ldots, \boldsymbol{c}_m) \in (R^d)^m$ and $\boldsymbol{y}' = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m) \in (R^d)^m$. We have $\rho_{s,(\boldsymbol{c}_1,\ldots,\boldsymbol{c}_m)}(L) = \prod_{i=1}^m \rho_{s,\boldsymbol{c}_i}(M)$. Therefore, the vector $\boldsymbol{y}'$ has distribution $D_{L,s,\boldsymbol{c}'}$. We have:

$$\sum_{i=1}^m z_i \cdot (\boldsymbol{y}_i - \boldsymbol{c}_i) = \begin{bmatrix} \sum_{i=1}^m z_i \cdot (y_i^{(1)} - c_i^{(1)}) \\ \vdots \\ \sum_{i=1}^m z_i \cdot (y_i^{(d)} - c_i^{(d)}) \end{bmatrix}.$$

The $((j-1)n+k)$-th coordinate of $\sum_{i=1}^m z_i \cdot (\boldsymbol{y}_i - \boldsymbol{c}_i)$ is the k-th coordinate of $\sum_{i=1}^m z_i \cdot (y_i^{(j)} - c_i^{(j)})$. Each coefficient of $z_i \cdot (y_i^{(j)} - c_i^{(j)})$ is the inner product between the corresponding coefficients of $y_i^{(j)} - c_i^{(j)}$ and a permutation of the coefficients of z_i (with some of them multiplied by -1). It suffices to study $\sum_{i=1}^m \langle y_i^{(j)} - c_i^{(j)}, z_i'\rangle$, for some (z_i')'s with $\|z_i'\| = \|z_i\|$. Let $\boldsymbol{e}_j$ be the j-th element of the standard basis of R^d (such that $e_{j,j} = (1,0,\ldots,0)$ and $e_{j,k} = (0,\ldots,0)$ for $k \neq j$). We have:

$$\sum_{i=1}^m \langle y_i^{(j)} - c_i^{(j)}, z_i'\rangle = \langle \boldsymbol{y}' - \boldsymbol{c}', (z_1'\boldsymbol{e}_j, \ldots, z_m'\boldsymbol{e}_j)\rangle = \|\boldsymbol{z}'\| \cdot \langle \boldsymbol{y}' - \boldsymbol{c}', \boldsymbol{w}_j\rangle,$$

where $\boldsymbol{w}_j \in K^{dm}$ is the unitary vector parallel to $(z_1'\boldsymbol{e}_j, \ldots, z_m'\boldsymbol{e}_j) \in K^{dm}$. By Lemma 7, we have, for all $j \in [d]$:

$$\Pr_{\boldsymbol{y}' \hookleftarrow D_{L,s,\boldsymbol{c}'}}[|\langle \boldsymbol{y}' - \boldsymbol{c}', \boldsymbol{w}_j\rangle| \geq st] \leq \frac{1+\varepsilon}{1-\varepsilon} t\sqrt{2\pi e} \cdot e^{-\pi t^2}.$$

The claim follows by applying the union bound over all $k \in [n]$ and all $j \in [d]$. □

[5] For the sake of simplicity, we identify elements x of K with their polynomial representations $\sigma_P(x)$.

The product of a continuous Gaussian on $\mathbb{R}$ with parameter s and a scalar $x \in \mathbb{R}$ is a continuous Gaussian with parameter xs. This can be generalized to the ring and module settings. The following result is given in [14], but without proof.

Lemma 9. *Let $\boldsymbol{s} \in \mathbb{R}^n$, $x \in K$ such that $\sigma_H(x)$ is a sample from $\nu_{\boldsymbol{s}}$ and $e \in K$ fixed. Then $\sigma_H(x \cdot e)$ is distributed from $\nu_{\boldsymbol{s'}}$ with, for all $1 \leq j \leq n/2$:*

$$s'_j = \sqrt{\mathrm{Re}(\sigma_j(e))^2 s_j^2 + \mathrm{Im}(\sigma_j(e))^2 s_{j+n/2}^2} \text{ and } s'_{j+n/2} = \sqrt{\mathrm{Im}(\sigma_j(e))^2 s_j^2 + \mathrm{Re}(\sigma_j(e))^2 s_{j+n/2}^2}.$$

In particular, if $s_j = s_{j+n/2}$ for all $1 \leq j \leq n/2$, then $s'_j = s_j |\sigma_j(e)|$.

Proof. Let us write $\sigma_C(x) = \sum_j x_j \cdot \boldsymbol{h}_j$ where each x_j is a sample of ν_{s_j}. By definition of the $\boldsymbol{h}_j$'s, we have $\sigma_j(x) = (x_j + i x_{j+n/2})$ and $\sigma_{j+n/2}(x) = (x_j - i x_{j+n/2})$, for $1 \leq j \leq n/2$. As a consequence:

$$\begin{aligned}\sigma_C(e \cdot x) =& (\sigma_1(e)(x_1 + i x_{n/2+1}), \ldots, \sigma_{n/2}(e)(x_{n/2} + i x_n), \\ & \sigma_{n/2+1}(e)(x_1 - i x_{n/2+1}), \ldots, \sigma_n(e)(x_{n/2} - i x_n)).\end{aligned}$$

For $1 \leq j \leq n/2$, we write $\sigma_j(e) = \alpha_j + i\beta_j$. This gives $\sigma_j(e \cdot x) = (\alpha_j x_j - \beta_j x_{j+n/2}) + i(\beta_j x_j + \alpha_j x_{j+n/2})$. We want to compute the y_j's such that $\sigma_C(e \cdot x) = \sum_j y_j \cdot \boldsymbol{h}_j$. For $1 \leq j \leq n/2$, we have:

$$y_j = \frac{1}{2}(\sigma_j(e \cdot x) + \sigma_{j+n/2}(e \cdot x)) \text{ and } y_{j+n/2} = \frac{i}{2}(\sigma_{j+n/2}(e \cdot x) - \sigma_j(e \cdot x)).$$

Therefore:

$$\begin{bmatrix} y_j \\ y_{j+n/2} \end{bmatrix} = \begin{bmatrix} \alpha_j & -\beta_j \\ \beta_j & \alpha_j \end{bmatrix} \begin{bmatrix} x_j \\ x_{j+n/2} \end{bmatrix},$$

where x_j is a sample from ν_{s_j}. The vector $(y_j, y_{j+n/2})$ is a full rank transformation of the vector $(x_j, x_{j+n/2})$, and thus y_j and $y_{j+n/2}$ are statically independent. Furthermore, the reals y_j and $y_{j+n/2}$ are samples of $\nu_{s'_j}$ and $\nu_{s'_{j+n/2}}$ respectively, with $s'_j = \sqrt{\alpha_j^2 s_j^2 + \beta_j^2 s_{j+n/2}^2}$ and $s'_{j+n/2} = \sqrt{\beta_j^2 s_j^2 + \alpha_j^2 s_{j+n/2}^2}$. □

Lemma 10. *Let $\boldsymbol{s} \in \mathbb{R}^n$, $\boldsymbol{x} \in K^d$ such that $\sigma_H(\boldsymbol{x})$ is chosen from $\nu_{\boldsymbol{s},\ldots,\boldsymbol{s}}$ and let $\boldsymbol{e} \in K^d$. Then $\sigma_H(\langle \boldsymbol{x}, \boldsymbol{e} \rangle)$ is distributed from $\nu_{\boldsymbol{s'}}$ with, for all $1 \leq j \leq n/2$:*

- $s'_j = \sqrt{\sum_{k=1}^d \left[\mathrm{Re}(\sigma_j(e_k))^2 s_j^2 + \mathrm{Im}(\sigma_j(e_k))^2 s_{j+n/2}^2\right]}$,
- $s'_{j+n/2} = \sqrt{\sum_{k=1}^d \left[\mathrm{Im}(\sigma_j(e_k))^2 s_j^2 + \mathrm{Re}(\sigma_j(e_k))^2 s_{j+n/2}^2\right]}$.

In particular, if $s_j = s_{j+n/2}$ for all $1 \leq j \leq n/2$, then $s'_j = s_j \cdot \sqrt{\sum_{k=1}^d |\sigma_j(e_k)|^2}$.

Proof. By Lemma 9, we have that $\sigma_H(x_k \cdot e_k)$ has distribution $\nu_{\boldsymbol{s}'_k}$ with for every $1 \leq j \leq n/2$:

$$s'_{k,j} = \sqrt{\mathrm{Re}(\sigma_j(e_k))^2 s_j^2 + \mathrm{Im}(\sigma_j(e_k))^2 s_{j+n/2}^2} \text{ and } s'_{k,j+n/2} = \sqrt{\mathrm{Im}(\sigma_j(e_k))^2 s_j^2 + \mathrm{Re}(\sigma_j(e_k))^2 s_{j+n/2}^2}.$$

The equality $\langle \boldsymbol{x}, \boldsymbol{e} \rangle = \sum_{k=1}^d x_k \cdot e_k$ completes the proof. □

3 The Short Integer Solution Problem

In this section, we describe a reduction from Mod-GIVP to M-SIS. To ease the presentation, we identify elements of K with their polynomial representations.

3.1 Variants of SIS

We first recall the SIS and R-SIS problems, and introduce M-SIS.

Definition 1. *The* Small Integer Solution problem $\mathrm{SIS}_{q,m,\beta}$ *is as follows: Given* $\boldsymbol{A} \in \mathbb{Z}_q^{n\times m}$ *chosen from the uniform distribution, find* $\boldsymbol{z} \in \mathbb{Z}^m$ *such that* $\boldsymbol{A}\boldsymbol{z} = 0 \bmod q$ *and* $0 < \|\boldsymbol{z}\| \leq \beta$.

As observed in [21, Le. 5.2], for any q, $\boldsymbol{A} \in \mathbb{Z}_q^{n\times m}$ and $\beta \geq \sqrt{m}q^{n/m}$, the SIS instance $(q, \boldsymbol{A}, \beta)$ admits a solution. There are several reductions from GIVP to SIS (see, e.g., [1,20,9]). The strongest known result is the following.

Theorem 2 (Adapted from [9, Th. 9.2]). *For* $\varepsilon(n) = n^{-\omega(1)}$ *(resp.* $\varepsilon(n) = 2^{-\Omega(n)}$*), there is a probabilistic polynomial time reduction from solving* $\mathrm{GIVP}_\gamma^{\eta_\varepsilon}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) to solving* $\mathrm{SIS}_{q,m,\beta}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability), for any* $m(n), q(n), \beta(n)$ *and* $\gamma(n)$ *such that* $\gamma \geq \beta\sqrt{n}\cdot\omega(\sqrt{\log n})$ *(resp.* $\gamma \geq \beta\Omega(n)$*),* $q \geq \beta\sqrt{n}\cdot\omega(\log n)$ *(resp.* $q \geq \beta\Omega(n^{3/2})$*) and* $m, \log q \leq \mathrm{poly}(n)$.

The R-SIS problem was introduced by [26] and [13].

Definition 2. *The problem* $\text{R-SIS}_{q,m,\beta}$ *is as follows: Given* $a_1, \ldots, a_m \in R_q$ *chosen independently from the uniform distribution, find* $z_1, \ldots, z_m \in R$ *such that* $\sum_{i=1}^m a_i \cdot z_i = 0 \bmod q$ *and* $0 < \|\boldsymbol{z}\| \leq \beta$*, where* $\boldsymbol{z} = (z_1|\ldots|z_m) \in \mathbb{Z}^{mn}$.

This problem over polynomials can be interpreted in terms of matrices. It is a variant of SIS where $\boldsymbol{A}$ is restricted to being block negacirculant: $\boldsymbol{A} = [\mathrm{Rot}(a_1)|\ldots|\mathrm{Rot}(a_m)]$, with:

$$\mathrm{Rot}(b) := \begin{bmatrix} b_0 & -b_{n-1} & \cdots & -b_1 \\ b_1 & b_0 & \cdots & -b_2 \\ \vdots & \vdots & \ddots & \vdots \\ b_{n-1} & b_{n-2} & \cdots & b_0 \end{bmatrix}, \text{ for } b = \sum_{i=0}^{n-1} b_i x^i \in R.$$

By using ideas from [14] (namely, the isomorphism between I/qI and R_q described in Subsection 2.1) into the proof of [13], one obtains the following result.

Theorem 3. *For* $\varepsilon(n) = n^{-\omega(1)}$ *(resp.* $\varepsilon(n) = 2^{-\Omega(n)}$*), there is a probabilistic polynomial time reduction from solving* $\mathrm{GIVP}_\gamma^{\eta_\varepsilon}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) to solving* $\text{R-SIS}_{q,m,\beta}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability), for any* $m(n), q(n), \beta(n)$ *and* $\gamma(n)$ *such that* $\gamma \geq \beta\sqrt{n}\cdot\omega(\sqrt{\log n})$ *(resp.* $\gamma \geq \beta\Omega(n)$*),* $q \geq \beta\sqrt{n}\cdot\omega(\log n)$ *(resp.* $q \geq \beta\Omega(n^{3/2})$*) and* $m, \log q \leq \mathrm{poly}(n)$.

The problem M-SIS generalizes both SIS and R-SIS.

Definition 3. *The problem* $\text{M-SIS}_{q,m,\beta}$ *is as follows: Given* $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_m \in (R_q)^d$ *chosen independently from the uniform distribution, find* $z_1, \ldots, z_m \in R$ *such that* $\sum_{i=1}^m \boldsymbol{a}_i \cdot z_i = 0 \bmod q$ *and* $0 < \|\boldsymbol{z}\| \leq \beta$*, where* $\boldsymbol{z} = (z_1|\ldots|z_m) \in \mathbb{Z}^{mn}$.

Like R-SIS, the M-SIS problems can be interpreted in terms of matrices. It consists in taking a SIS matrix $\boldsymbol{A}$ of the form:

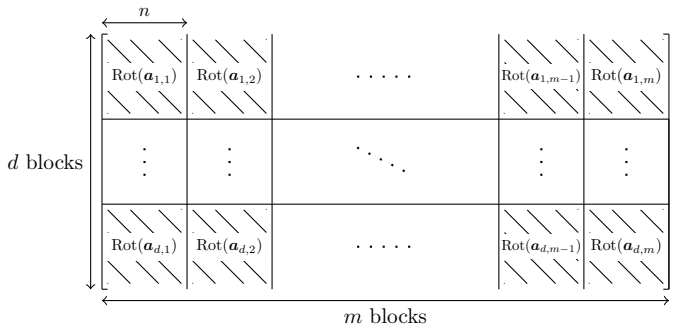

In the rest of this section, we will prove the following result.

Theorem 4. *For any $d \geq 1$ and $\varepsilon(nd) = (nd)^{-\omega(1)}$ (resp. $\varepsilon(nd) = 2^{-\Omega(nd)}$), there is a probabilistic polynomial time reduction from solving* Mod-GIVP$^{\eta_\varepsilon}_{\gamma}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) to solving* M-SIS$_{q,m,\beta}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability), for any $m(nd), q(nd), \beta(dn)$ and $\gamma(nd)$ such that $\gamma \geq \beta\sqrt{nd}\cdot\omega(\sqrt{\log nd})$ (resp. $\gamma \geq \beta\Omega(nd)$), $q \geq \beta\sqrt{nd}\cdot\omega(\log nd)$ (resp. $q \geq \beta\Omega((nd)^{3/2})$) and $m, \log q \leq \mathrm{poly}(nd)$.*

Taking $d = 1$ in Theorem 4 allows us to recover Theorem 3. Also, by taking $d = n$ and $n = 1$ in Theorem 4, we obtain a hardness result for SIS that is as good as that of Theorem 2.

3.2 A reduction from Mod-GIVP to M-SIS

In order to prove that the new problem M-SIS is as hard as GIVP restricted to module lattices, we use the following intermediate problem, introduced in [20].

Definition 4 ([20, Def. 5.3]). *The* Incremental Independent Vectors Problem IncGIVP$^{\eta_\varepsilon}_{\gamma}$*, is as follows: Given a tuple $(\boldsymbol{B}, \boldsymbol{S}, \mathcal{H})$ where $\boldsymbol{B}$ is a basis of an n-dimensional lattice, $\boldsymbol{S} \subseteq \mathcal{L}(\boldsymbol{B})$ is a full-rank set of vectors such that $\|\boldsymbol{S}\| \geq \gamma \cdot \eta_\varepsilon(\mathcal{L}(\boldsymbol{B}))$ and $\mathcal{H}$ is a hyperplane, find $\boldsymbol{h} \in \mathcal{L}(\boldsymbol{B}) \setminus \mathcal{H}$ such that $\|\boldsymbol{h}\| \leq \|\boldsymbol{S}\|/2$.*

Theorem 5 ([16, Th. 6.3]). *For any function ε and γ, there is a probabilistic polynomial time reduction from solving* GIVP$^{\eta_\varepsilon}_{\gamma}$ *(in the worst case, with high probability) to solving* IncGIVP$^{\eta_\varepsilon}_{\gamma}$ *(in the worst case, with high probability).*

As the latter reduction preserves the lattice, it induces a reduction from Mod-GIVP$^{\eta_\varepsilon}_{\gamma}$ to Mod-IncGIVP$^{\eta_\varepsilon}_{\gamma}$, i.e., IncGIVP$^{\eta_\varepsilon}_{\gamma}$ restricted to module lattices. To prove Theorem 4, we provide a reduction from Mod-IncGIVP$^{\eta_\varepsilon}_{\gamma}$ to M-SIS$_{q,m,\beta}$.

Suppose that an oracle $\mathcal{O}$ solves M-SIS$_{q,m,\beta}$ with probability $(nd)^{-O(1)}$ (resp. $2^{-o(nd)}$). The algorithm for Mod-IncGIVP proceeds as follows on input $(\boldsymbol{B}, \boldsymbol{S}, \mathcal{H})$. We write $M = \mathcal{L}(\mathbf{B})$. Let s be such that

$$\max\left(\frac{2q}{\gamma}, \omega(\sqrt{\log nd})\right)\|\boldsymbol{S}\| \leq s \leq \frac{q\|\boldsymbol{S}\|}{2\beta\sqrt{nd}\cdot\omega(\sqrt{\log nd})} \quad (\text{resp. } \max(\frac{2q}{\gamma}, \Omega(\sqrt{nd}))\|\boldsymbol{S}\| \leq s \leq \frac{q\|\boldsymbol{S}\|}{2\beta\Omega(nd)}).$$

- For $i \leq m$, let $\boldsymbol{y}_i$ be sampled from $D_{\mathcal{L}(\mathbf{B}),s,0}$ (using Theorem 1), and $\boldsymbol{a}_i = \Theta^{-1}(\boldsymbol{y}_i)$ (see Section 2.1).
- Invoke oracle $\mathcal{O}$ on input $(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_m)$. If $\mathcal{O}$ succeeds, it returns $\boldsymbol{z} = (z_1|\ldots|z_m) \in R^m$ such that $\sum_{i=1}^m \boldsymbol{a}_i \cdot z_i = 0 \bmod q$ and $0 < \|\boldsymbol{z}\| \leq \beta$.
- Output $\boldsymbol{h} = \frac{1}{q}\sum_{i=1}^m z_i \cdot \boldsymbol{y}_i$.

This algorithm runs in polynomial time (without considering the run-time of oracle $\mathcal{O}$). Also, thanks to the parameter constraints, the interval to which the standard deviation s must belong is non-empty. Moreover, the standard deviation s is sufficiently large for the assumptions of Theorem 1 to hold. Indeed, by [18, Le. 7.1] and given M and $\mathbf{S}$, it is possible to compute (in polynomial time) a basis $\mathbf{T}$ of M such that $\|\tilde{\mathbf{T}}\| \leq \|\tilde{\mathbf{S}}\| \leq \|\mathbf{S}\|$. We use this basis and we have that $s \geq \|\tilde{\mathbf{T}}\| \cdot \omega(\sqrt{\log nd})$ (resp. $s \geq \|\tilde{\mathbf{T}}\| \cdot \Omega(\sqrt{nd})$).

Lemma 11. *The statistical distance between the distribution of $(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_m)$ and the uniform distribution over $(R_q)^d$ is a most 2ε.*

Proof. We have $s \geq \frac{2q}{\gamma} \cdot \|\boldsymbol{S}\|$ and $\|\boldsymbol{S}\| \geq \gamma \cdot \eta_\varepsilon(M)$. This implies that $s \geq q \cdot \eta_\varepsilon(M) = \eta_\varepsilon(qM)$. By Lemma 4 applied to the lattices M and qM, the statistical distance between the distribution of $(\boldsymbol{y}_i \bmod qM)$ and the uniform distribution on M/qM is at most 2ε. As Θ^{-1} is an isomorphism from M/qM to $(R/qR)^d$, the statistical distance between the distribution of the $\boldsymbol{a}_i = \Theta^{-1}(\boldsymbol{y}_i)$ and the uniform distribution on $(R/qR)^d$ is also at most 2ε. □

As a consequence, the oracle $\mathcal{O}$ succeeds with probability $(nd)^{-O(1)}$ (resp. $2^{-o(nd)}$). In the following, we assume we are in that situation.

Lemma 12. *For any hyperplane $\mathcal{H}$, the probability that the output vector $\boldsymbol{h}$ does not belong to $\mathcal{H}$ is $\geq 1/100$.*

Proof. As $\mathcal{O}$ succeeded, the vector $\boldsymbol{z}$ is non-zero. By definition of $\boldsymbol{h}$, for every $\boldsymbol{y}'_1$ we have:

$$\begin{aligned}\boldsymbol{h} \in \mathcal{H} &\Leftrightarrow \sum_{i=1}^{m} z_i \cdot \boldsymbol{y}_i \in \mathcal{H} \Leftrightarrow \ z_1 \cdot \boldsymbol{y}_1 \in -\sum_{i=2}^{m} z_i \cdot \boldsymbol{y}_i + \mathcal{H} \\ &\Leftrightarrow (\boldsymbol{y}_1 - \boldsymbol{y}'_1) \in -\boldsymbol{y}'_1 + \frac{1}{z_1}(\mathcal{H} - \sum_{i=2}^{m} z_i \cdot \boldsymbol{y}_i) = \mathcal{H}'.\end{aligned}$$

Assume that we fix $\boldsymbol{y}'_1 = \boldsymbol{y}_1 \bmod qM$, then $\boldsymbol{y}_1 = \boldsymbol{y}'_1 + \boldsymbol{y}''_1$, with $\boldsymbol{y}'_1$ fixed and the vector $\boldsymbol{y}''_1$ statistically independent of all the $\boldsymbol{a}_i$'s, z_i's and $\boldsymbol{y}_i$'s for $i > 1$. The conditional distribution of $\boldsymbol{y}''_1 = (\boldsymbol{y}_1 - \boldsymbol{y}'_1)$ is $D_{qM,s,-\boldsymbol{y}'_1}$. Therefore:

$$\Pr\left[(\boldsymbol{y}_1 - \boldsymbol{y}'_1) \notin \mathcal{H}' | \boldsymbol{y}'_1, (\boldsymbol{a}_1, \ldots, \boldsymbol{a}_m), (z_1, \ldots, z_m)\right] = \Pr_{\boldsymbol{y}''_1 \sim D_{qM,s,-\boldsymbol{y}'_1}}[\boldsymbol{y}''_1 \notin \mathcal{H}'].$$

As $s \geq 2q \cdot \eta_\varepsilon(M) = 2\eta_\varepsilon(qM)$, Lemma 5 gives that this probability is $\geq 1/100$. □

The following completes the proof of Theorem 4.

Lemma 13. *We have $\boldsymbol{h} \in M$ and, with probability close to 1, we have that $\|\boldsymbol{h}\| \leq \|\boldsymbol{S}\|/2$.*

Proof. Let us first show that $\boldsymbol{h} \in M$. We have :

$$\boldsymbol{h} = \frac{1}{q}\sum_{i=1}^{m} z_i \cdot \boldsymbol{y}_i = \frac{1}{q}\sum_{i=1}^{m} z_i \cdot \Theta(\boldsymbol{a}_i) = \frac{1}{q}\sum_{i=1}^{m} z_i \cdot \left(\sum_{j=1}^{d}(t_j a_{i,j} + q s_{i,j}) \cdot \boldsymbol{b}_j\right),$$

where $s_{i,j} \in I_j$ and $t_j \in I_j$. Therefore:

$$\boldsymbol{h} = \frac{1}{q}\sum_{j=1}^{d} t_j \left(\sum_{i=1}^{m} z_i \cdot a_{i,j}\right) \cdot \boldsymbol{b}_j + \sum_{j=1}^{d}\left(\sum_{i=1}^{m} z_i \cdot s_{i,j}\right) \cdot \boldsymbol{b}_j.$$

As $s_{i,j} \in I_j$ for all j, we have $\sum_{j=1}^{d}(\sum_{i=1}^{m} z_i \cdot s_{i,j}) \cdot \boldsymbol{b}_j \in M$. We also have $\sum_{i=1}^{m} z_i \cdot \boldsymbol{a}_i = 0 \bmod q$ and, as $t_j \in I_j$ for all j, we have that $t_j \cdot \frac{1}{q}(\sum_{i=1}^{m} z_i \cdot a_{i,j}) \in I_j$. This shows that $\boldsymbol{h} \in M$.

We now show that $\|\boldsymbol{h}\| \leq \|\boldsymbol{S}\|/2$. Recall that $\boldsymbol{h} = \frac{1}{q}\sum_{i=1}^{m} z_i \cdot \boldsymbol{y}_i$, then we have $\|\boldsymbol{h}\| = \frac{1}{q}\|\sum_{i=1}^{m} z_i \cdot \boldsymbol{y}_i\|$. As in the previous proof, we define $\boldsymbol{y}'_i = \boldsymbol{y}_i \bmod qM$. Then, we have $\boldsymbol{y}_i = \boldsymbol{y}''_i + \boldsymbol{y}'_i$ with $\boldsymbol{y}''_i$ statistically independent from the z_i's and distributed as $D_{qM,s,-\boldsymbol{y}'_i}$. By Lemma 8, for $s \geq \eta_\varepsilon(qM)$ and $t = \omega(\sqrt{\log nd})$ (resp. $t = \Omega(\sqrt{nd})$), we know that:

$$\Pr_{\boldsymbol{y}''_i \sim D_{qM,s,-\boldsymbol{y}'_i}}\left[\left\|\sum_{i=1}^{m} z_i \cdot (\boldsymbol{y}''_i + \boldsymbol{y}'_i)\right\| \geq st\sqrt{nd} \cdot \|\boldsymbol{z}\|\right] \leq (nd)^{-\omega(1)} \quad (\text{resp. } \leq 2^{-o(nd)}).$$

So, with probability close to 1, we have $\|\sum_{i=1}^{m} z_i \cdot \boldsymbol{y}_i\| \leq st\sqrt{nd} \cdot \|\boldsymbol{z}\|$. As $0 < \|\boldsymbol{z}\| \leq \beta$, we have:

$$\|\boldsymbol{h}\| = \frac{1}{q}\left\|\sum_{i=1}^{m} z_i \cdot \boldsymbol{y}_i\right\| \leq \frac{st\beta\sqrt{nd}}{q}.$$

Finally, since $s \leq \frac{q \cdot \|\boldsymbol{S}\|}{2\beta t\sqrt{nd}}$, we obtain $\|\boldsymbol{h}\| \leq \frac{\|\boldsymbol{S}\|}{2}$. □

4 Learning with errors over modules

In this section, we describe a reduction from Mod-GIVP to M-LWE (Learning With Errors over modules). To ease the presentation, we now identify elements of K with their σ_H embeddings.

4.1 Learning With Errors

We let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ denote the segment $[0, 1)$ with addition modulo 1.

Let us recall the following definitions from [31]. For a probability density function χ on $\mathbb{T}$ and a vector $\boldsymbol{s} \in \mathbb{Z}_q^n$, we let $A_{\boldsymbol{s},\chi}$ denote the distribution on $\mathbb{Z}_q^n \times \mathbb{T}$ obtained by choosing a vector $\boldsymbol{a} \in \mathbb{Z}_q^n$ uniformly at random, choosing $e \in \mathbb{T}$ according to χ, and returning $(\boldsymbol{a}, \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s}\rangle + e)$.

Definition 5. *The* search *version of the* Learning With Error problem $\mathrm{CLWE}_{q,\chi}$ *is as follows: Let* $\boldsymbol{s} \in \mathbb{Z}_q^n$ *be secret; Given arbitrarily many samples from* $A_{\boldsymbol{s},\chi}$*, the goal is to find* $\boldsymbol{s}$*.*

The decision *version of the* Learning With Error problem $\mathrm{DLWE}_{q,\chi}$ *is as follows: Let* $\boldsymbol{s} \in \mathbb{Z}_q^n$ *uniformly random; The goal is to distinguish between arbitrarily many independent samples from* $A_{s,\chi}$ *and the same number of independent samples from* $U(\mathbb{Z}_q^n \times \mathbb{T})$*.*

It is also possible to interpret LWE in terms of linear algebra: Suppose the number of requested samples $(\boldsymbol{a}_i, \frac{1}{q}\langle \boldsymbol{a}_i, \boldsymbol{s}\rangle + e_i)$ from $A_{s,\chi}$ is m, then we consider the matrix $\boldsymbol{A} \in \mathbb{Z}_q^{m \times n}$ whose rows are the $\boldsymbol{a}_i$'s, and we create the vector $\boldsymbol{e} = (e_1, \ldots, e_m)^T$. Then CLWE is as follows:

$$\overset{m}{\uparrow}\begin{bmatrix} \\ \boldsymbol{A} \\ \\ \end{bmatrix}\underset{n}{\longleftrightarrow}, \quad \frac{1}{q} \cdot \begin{bmatrix} \\ \boldsymbol{A} \\ \\ \end{bmatrix}\begin{bmatrix} s_1 \\ \vdots \\ s_n \end{bmatrix} + \begin{bmatrix} e_1 \\ \vdots \\ e_m \end{bmatrix} \xrightarrow{\text{find}} \boldsymbol{s}$$

Theorem 6 ([31]). *Let* $\varepsilon(n) = n^{-\omega(1)}$ *(resp.* $\varepsilon(n) = 2^{-\Omega(n)}$*),* $\alpha \in (0, 1)$ *and* $q \geq 2$ *such that* $\alpha q > 2\sqrt{n}$*. There exists a quantum reduction from solving* $\mathrm{GIVP}^{\eta_\varepsilon}_{\sqrt{8}n/\alpha}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) to solving* $\mathrm{CLWE}_{q,\nu_\alpha}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability).*

Assume that q *is prime,* $q \leq poly(n)$*, and that* χ *is a probability density function on* $\mathbb{T}$*. There exists a reduction from* $\mathrm{CLWE}_{q,\chi}$ *to* $\mathrm{DLWE}_{q,\chi}$*.*

The first main result of [31] is the reduction from GIVP to the computational version of LWE. It makes use of the following intermediary problem, where ϕ denotes an arbitrary real-valued function on lattice, called *Discrete Gaussian Sampling problem* (DGS_ϕ): Given an n-dimensional lattice Λ and a number $r > \phi(\Lambda)$, output a sample from $D_{\Lambda,r}$. Regev's reduction proceeds in two steps:

$$\mathrm{GIVP}_{\frac{\sqrt{8}n}{\alpha}} \xrightarrow{\text{[31, Le. 3.17]}} \mathrm{DGS}_{\frac{\sqrt{2n}}{\alpha}} \xrightarrow{\text{Lemma 14}} \mathrm{CLWE}_{q,\nu_\alpha}$$

The first reduction is lattice-preserving and also works for the structured versions of LWE to be considered later. Oppositely, the second one will need to be modified. It comes from the following result:

Lemma 14 ([31, Le. 3.3]). *Let $\varepsilon(n) = n^{-\omega(1)}$ (resp. $\varepsilon(n) = 2^{-\Omega(n)}$), $\alpha \in (0,1)$ and $q \geq 2$. Assume that we have access to an oracle that solves $\mathrm{CLWE}_{q,\nu_\alpha}$ in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability). Then there exists a polynomial time (resp. sub-exponential time) quantum algorithm that, given an n-dimensional lattice Λ, a number $r > \sqrt{2}q \cdot \eta_\varepsilon(\Lambda)$ and $\mathrm{poly}(n)$ (resp. $2^{o(n)}$) samples from $D_{\Lambda,r}$, produces a sample from $D_{\Lambda,\frac{r\sqrt{n}}{\alpha q}}$ with non-negligible (resp. non-exponentially small) probability.*

The principle of the Regev's reduction from DGS to CLWE is to use Lemma 14 several times to progressively decrease the value of r. Take $r > \sqrt{2}q \cdot \eta_\varepsilon(\Lambda)$ and $r_i = r \cdot (\alpha q/\sqrt{n})^i$. The first iteration starts with $r_{3n} > 2^{3n} > 2^{2n}\lambda_n(\Lambda)$ (using a LLL-reduction algorithm beforehand). Then it obtains $\mathrm{poly}(n)$ (resp. $2^{o(n)}$) samples of $D_{\Lambda,r_{3n}}$ by Theorem 1, and finishes with $\mathrm{poly}(n)$ (resp. $2^{o(n)}$) samples of $D_{\Lambda,r_{3n-1}}$ (the reduction repeats $\mathrm{poly}(n)$ (resp. $2^{o(n)}$) times the same iteration with the same samples in input to obtain sufficiently many different samples in output). It iterates until having $\mathrm{poly}(n)$ (resp. $2^{o(n)}$) samples of D_{Λ,r_1} with $r_1 = r\alpha q/\sqrt{n} > \sqrt{2}q \cdot \eta_\varepsilon(\Lambda)$ then it iterates a last time to obtain samples of D_{Λ,r_0} with $r_0 = r > \sqrt{2n} \cdot \eta_\varepsilon(\Lambda)/\alpha$. These samples are solutions to $\mathrm{DGS}_{\sqrt{2n}\cdot\eta_\varepsilon(\Lambda)/\alpha}$.

To prove Lemma 14, Regev uses the intermediary problems called q-BDD_δ: Given a lattice Λ and any point $\boldsymbol{y} \in \mathbb{R}^n$ within distance $\delta < \lambda_1(\Lambda)/2$ of the lattice, output the coset of $\Lambda/q\Lambda$ of the closest vector to $\boldsymbol{y}$. The proof of Lemma 14 consists also of a sequence of reductions:

$$\mathrm{DGS}_{\Lambda,\frac{r\sqrt{n}}{\alpha q}} \xrightarrow[\text{(quantum)}]{\text{[31, Le. 3.14 \& 3.5]}} q\text{-}\mathrm{BDD}_{\Lambda^*,\frac{\alpha q}{\sqrt{2}r}} \xrightarrow{\text{Lemma 15}} \begin{array}{c}\mathrm{CLWE}_{q,\nu_\alpha}\\ +\\ \text{samples from } D_{\Lambda,r}\end{array}$$

The first reduction also works for the structured versions of LWE to be considered later. However, we will modify the second reduction, by proving an adaptation of the following result.

Lemma 15 ([31, Le. 3.4]). *Let $\varepsilon(n) = n^{-\omega(1)}$ (resp. $\varepsilon(n) = 2^{-\Omega(n)}$), $\alpha \in (0,1)$ and $q \geq 2$. Let Λ be a n-dimensional lattice and $r \geq \sqrt{2}q \cdot \eta_\varepsilon(\Lambda)$. Given access to an oracle sampling from the distribution $D_{I,r}$, there exists a probabilistic reduction from solving q-$\mathrm{BDD}_{\Lambda^*,\frac{\alpha q}{\sqrt{2}r}}$ in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) to solving $\mathrm{CLWE}_{q,\nu_\alpha}$ in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability).*

Regev's reduction from CLWE *to* DLWE. The second main result from [31] is a reduction from the computational problem CLWE to its decisional counterpart DLWE. This reduction does not carry over to the structured variants of LWE.

4.2 Learning With Errors over rings

The R-LWE problem was introduced by Lyubashevsky et al in [14]. Let ψ be a distribution on $\mathbb{T}_{R^*} = K_{\mathbb{R}}/R^*$ and $s \in R_q^*$. (Recall that for our choice of K, we have $R^* = \frac{1}{n}R$.) We let $A_{s,\psi}^{(R)}$ denote the distribution on $R_q \times \mathbb{T}_{R^*}$ obtained by choosing $a \in R_q$ uniformly at random and $e \in \mathbb{T}_{R^*}$ according to ψ, and returning $(a, (a \cdot s)/q + e)$.

We also recall the distribution Υ_α used in [14]. The gamma distribution $\Gamma(2,1)$ with shape parameter 2 and scale parameter 1 has density $x\exp(-x)$ for $x \geq 0$ and zero for $x < 0$. For $\alpha > 0$, a distribution sampled from Υ_α is an elliptical Gaussian distribution $\nu_{\boldsymbol{r}}$ whose parameters are $r_i = r_{i+n/2} = \alpha\sqrt{1 + \sqrt{n}x_i}$, where $x_1, \ldots, x_{n/2}$ are chosen independently from $\Gamma(2,1)$.

Definition 6. *Let $q \geq 2$ and Ψ be a family of distributions on $\mathbb{T}_{R^*}$. The* search *version of the* Ring Learning With Error problem R-CLWE$_{q,\Psi}$ *is as follows: Let $s \in R_q^*$ be secret and $\psi \in \Psi$; Given arbitrarily many samples from $A_{s,\psi}^{(R)}$, the goal is to find s.*

Let Υ be a distribution over a family of noise distributions over $K_{\mathbb{R}}$. The decision *version of the* Ring Learning With Error problem R-DLWE$_{q,\Upsilon}$ *is as follows: Let $s \in R_q^*$ uniformly random and ψ sampled from Υ; The goal is to distinguish between arbitrarily many independent samples from $A_{\boldsymbol{s},\psi}^{(R)}$ and the same number of independent samples from $U(R_q, \mathbb{T}_{R^*})$.*

As for R-SIS, this problem can be interpreted in terms of linear algebra. It is a variant of LWE where the matrix $\boldsymbol{A}$ is restricted to being block-negacirculant: $\boldsymbol{A} = [\mathrm{Rot}(a_1)|\dots|\mathrm{Rot}(a_m)]^T$. The two main results from [14] are a reduction from Id-GIVP to R-CLWE and a reduction from the search version R-CLWE to the decision version R-DLWE.

Theorem 7 ([14, Th. 4.1 & Th. 5.2]). *Let $\varepsilon(n) = n^{-\omega(1)}$ (resp. $\varepsilon(n) = 2^{-\Omega(n)}$), $\alpha \in (0,1)$ and $q \geq 2$ of known factorization such that $\alpha q > \omega(\sqrt{\log n})$ (resp. $\Omega(\sqrt{n})$). There exists a quantum reduction from solving* Id-GIVP$_{\gamma}^{\eta_\varepsilon}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) to solving* R-CLWE$_{q,\Psi_{\leq\alpha}}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) with $\gamma = 2\sqrt{n} \cdot \omega(\sqrt{\log n})/\alpha$ (resp. $2\Omega(n)/\alpha$).*

Assume that q is prime, $q \leq \mathrm{poly}(n)$, and that $x^n + 1$ has n linear factors modulo q (i.e., we have $q = 1 \bmod 2n$). There exists a randomized reduction from R-CLWE$_{q,\Psi_{\leq\alpha}}$ *to* R-DLWE$_{q,\Upsilon_\alpha}$.

The Lyubashevsky et al reduction from Id-GIVP to R-CLWE relies on the same sequence of reductions as Regev's proof of hardness of CLWE, but with problems restricted to ideal lattices. The only step in Regev's reduction that fails to carry over to the ideal/ring setting is Lemma 15. Lyubashevsky et al circumvent it by proving the following. In this Lemma, the problem q-Id-BDD is the restriction of q-BDD to ideal lattice lattices and instead of using the Euclidean norm for bounding the distance to the lattice, they use the infinity norm.

Lemma 16 ([14, Le. 4.4]). *Let $\varepsilon = n^{-\omega(1)}$ (resp. $\varepsilon(n) = 2^{-\Omega(n)}$), $\alpha \in (0,1)$ and $q \geq 2$ of known factorization. Let $I \subseteq R$ be an ideal and $r \geq \sqrt{2}q \cdot \eta_\varepsilon(I)$. Given access to an oracle sampling from the distribution $D_{I,r}$, there exists a probabilistic reduction from solving q-Id-BDD$_{I^*,\frac{\alpha q}{\sqrt{2}r}}$ in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) to solving* R-CLWE$_{q,\Psi_{\leq\alpha}}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability).*

The reduction from R-CLWE to R-DLWE from [14, Th. 5.2] proceeds by several reductions between intermediates problems, which we will also consider in our reduction for the module variant of LWE. Let $q = 1 \bmod 2n$ be prime, then $(q) = \prod_{i=1}^n \mathfrak{q}_i$ where any $\mathfrak{q}_i$ is a prime ideal with norm $N(\mathfrak{q}_i) = q$. Lyubashevsky et al define:

- $\mathfrak{q}_i$**-RLWE**$_{q,\Psi}$, with parameters Ψ a family of distributions over $\mathbb{T}_{R^*}$ and $i \leq n$: Given access to an oracle sampling from $A_{s,\psi}^{(R)}$ for an arbitrary $s \in R_q^*$ and $\psi \in \Psi$, find $s \bmod \mathfrak{q}_i R_q^*$.
- **Hybrid distribution** $A_{s,\psi}^{(R,i)}$, with parameters ψ a distribution over $\mathbb{T}_{R^*}$, $s \in R_q^*$, and $i \leq n$: The distribution $A_{s,\psi}^{(R,i)}$ over $R_q \times \mathbb{T}_{R^*}$ is defined as follows: Choose (a,b) from $A_{s,\psi}^{(R)}$ and return $(a, b + r/q)$ where r is uniformly random and independent in $R_q^*/\mathfrak{q}_j R^*$ for all $j \leq i$, and is 0 modulo the remaining $\mathfrak{q}_j R^*$'s.
- **DecRLWE**$_{q,\Psi}^i$, with parameters Ψ a family of distributions on $\mathbb{T}_{R^*}$ and $i \leq n$: Given access to an oracle sampling from $A_{s,\psi}^{(R,j)}$ for an arbitrary $s \in R_q^*$, $\psi \in \Psi$ and $j \in \{i-, i\}$, find j.

The sequence of reductions is as follows:

$$\text{R-CLWE}_{q,\Psi} \xrightarrow{\text{[14, Le. 5.5]}} \mathfrak{q}_i\text{-RLWE}_{q,\Psi} \xrightarrow{\text{[14, Le. 5.8]}} \text{DecRLWE}^i_{q,\Psi} \xrightarrow{\text{[14, Le. 5.11 \& 5.13]}} \text{R-DLWE}_{q,\Upsilon}$$

In our adaptation to modules, we will keep the general structure of this reduction. The two first intermediates reductions will be modified, while the reduction from DecRLWE to R-DLWE will be kept as it also works in the case of modules.

4.3 Learning With Errors over modules

The M-LWE problem generalizes both LWE and R-LWE, and was recently introduced in [6]. Let ψ be some probability distribution on $\mathbb{T}_{R^*}$ and $\boldsymbol{s} \in (R_q^*)^d$ be a vector. We define $A^{(M)}_{s,\psi}$ as the distribution on $(R_q)^d \times \mathbb{T}_{R^*}$ obtained by choosing a vector $\boldsymbol{a} \in (R_q)^d$ uniformly at random, and $e \in \mathbb{T}_{R^*}$ according to ψ, and returning $(\boldsymbol{a}, \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s}\rangle + e)$.

Definition 7. *Let $q \geq 2$ and Ψ be a family of distributions on $\mathbb{T}_{R^*}$. The* search *version of the* Module Learning With Error problem M-CLWE$_{q,\Psi}$ *is as follows: Let $\boldsymbol{s} \in (R_q^*)^d$ be secret and $\psi \in \Psi$; Given arbitrarily many samples from $A^{(M)}_{\boldsymbol{s},\psi}$, the goal is to find $\boldsymbol{s}$.*

For an integer $q \geq 2$ and a distribution Υ over a family of distributions over $K_{\mathbb{R}}$. The decision *version of the* Module Learning With Error problem M-DLWE$_{q,\Upsilon}$ *is as follows: Let $\boldsymbol{s} \in (R_q^*)^d$ uniformly random and ψ sampled from Υ; The goal is to distinguish between arbitrarily many independent samples from $A^{(M)}_{\boldsymbol{s},\psi}$ and the same number of independent samples from $U((R_q)^d, \mathbb{T}_{R^*})$.*

As for LWE and R-LWE, the problem M-LWE can be interpreted in terms of linear algebra. It consists in taking the LWE matrix $\boldsymbol{A}$ of the form:

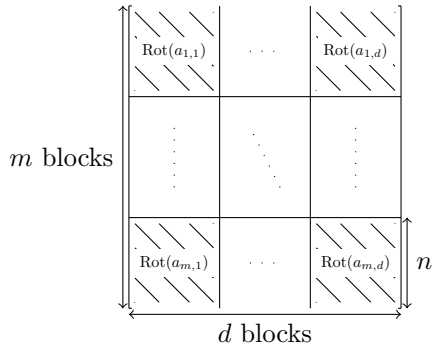

Theorem 8. *Let $\varepsilon(n) = n^{-\omega(1)}$ (resp. $\varepsilon(n) = 2^{-\Omega(n)}$), $\alpha \in (0,1)$ and $q \geq 2$ of known factorization such that $\alpha q > 2\sqrt{d} \cdot \omega(\sqrt{\log n})$ (resp. $\alpha q > 2\sqrt{d} \cdot \Omega(\sqrt{n})$). There is a quantum reduction from solving* Mod-GIVP$^{\eta_\varepsilon}_{\gamma}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) to solving* M-CLWE$_{q,\Psi_{\leq\alpha}}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) with $\gamma = \sqrt{8n}d \cdot \omega(\sqrt{\log n})/\alpha$ (resp. $d \cdot \Omega(n)/\alpha$).*

Assume that q is prime, $q \leq \text{poly}(n)$ and that $x^n + 1$ has n linear factors modulo q (i.e., we have $q = 1 \bmod 2n$). There exists a probabilistic reduction from M-CLWE$_{q,\Psi_{\leq\alpha}}$ *to* M-DLWE$_{q,\Upsilon_\alpha}$.

Notice that by taking $d = 1$ in Theorem 8 we recover Theorem 7, and that by taking $d = n$ and $n = 1$ we obtain a slighlty weaker variant of Theorem 6 (because of the use of Υ_α rather than ν_α). The remaining of this section is devoted to proving Theorem 8.

A reduction from Mod-GIVP to M-CLWE. The reduction from Mod-GIVP to M-CLWE follows the same design principle as Regev's reduction from GIVP to CLWE. The only component we modify is the reduction from q-BDD to CLWE that is given access to samples from $D_{L,r}$. More precisely, we replace Lemma 15 by Lemma 17 below. This proves the first part of Theorem 8.

We define the problem q-Mod-BDD as the restriction of q-BDD to module lattices, with the following variation. Instead of using the Euclidean norm for bounding the distance to the lattice, we use $\|\cdot\|_{2,\infty}$, defined by $\|\boldsymbol{e}\|_{2,\infty} = \max_i \sqrt{\sum_{k=1}^{d} |\sigma_i(e_k)|^2}$ for $\boldsymbol{e} \in K^d$.

Lemma 17. *Let $\varepsilon = n^{-\omega(1)}$ (resp. $\varepsilon(n) = 2^{-\Omega(n)}$), $\alpha \in (0,1)$ and $q \geq 2$. Let $M \subseteq R^d$ be an R-module, and $r > \sqrt{2}q \cdot \eta_\varepsilon(M)$. Given access to an oracle sampling from the distribution $D_{M,r}$, there exists a probabilistic reduction from solving q-Mod-BDD$_{M^*, \frac{\alpha q}{\sqrt{2}r}}$ in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability) to solving* M-CLWE$_{q,\Psi_{\leq\alpha}}$ *in polynomial time with non-negligible probability (resp. in sub-exponential time with non-exponentially small probability).*

The principle of the reduction is to construct from $\boldsymbol{y}$, the input of q-Mod-BDD, and from some discrete and continuous Gaussian samples, the pairs $(\boldsymbol{a}, b)$ distributed as $A^{(M)}_{\boldsymbol{s},\psi}$, where $\boldsymbol{s}$ will directly depend on the closest vector $\boldsymbol{x}$ to $\boldsymbol{y}$. To produce such samples $(\boldsymbol{a}, b)$ with the desired distribution, we combine the corresponding proofs for LWE and R-LWE (those of Lemmata 15 and 16). Then a call to the oracle of M-CLWE returns $\boldsymbol{s}$ and let us recover information on $\boldsymbol{x}$.

Proof of Lemma 17. Let $\mathcal{O}$ be the oracle which, given $m \leq poly(n)$ samples $(\boldsymbol{a}, b)$ from $A^{(M)}_{\boldsymbol{s},\psi}$ for $\psi \in \Psi_{\leq\alpha}$, outputs $\boldsymbol{s}$ in polynomial time (respectively subexponential time) with probability $(nd)^{-O(1)}$ (resp. $2^{-o(nd)}$). Given $M = \sum_{i=1}^{d} I_i \cdot \boldsymbol{b}_i$, the input of the reduction is $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{e}$ such that $\boldsymbol{x} \in M^*$ and $\|\boldsymbol{e}\|_{2,\infty} \leq \delta = \frac{\alpha q}{\sqrt{2}r}$. The goal is to find $\boldsymbol{x} \bmod qM^*$. The reduction is as follows:

- For all i, compute $t_i \in I_i$ such that $t_i \cdot I_i$ and (q) are coprime, and let $\boldsymbol{t} = (t_i)_{1 \leq i \leq d}$.
- To create an M-LWE sample:
  - Get a fresh $\boldsymbol{z}$ distributed as $D_{M,r}$ and a fresh e' distributed as $\nu_{\alpha/\sqrt{2}}$,
  - Let $\boldsymbol{a} = \Theta^{-1}(\boldsymbol{z} \bmod qM)$ and $b = \frac{1}{q}\langle \boldsymbol{z}, \boldsymbol{y} \rangle + e' \bmod R^*$ (see the definition of Θ in Section 2.1).
- Invoke $\mathcal{O}$ on input m samples $(\boldsymbol{a}, b)$ of M-LWE. If $\mathcal{O}$ succeeds, then it outputs some $\boldsymbol{s} \in (R_q^*)^d$.
- Return $\Theta^{-1}(\boldsymbol{s}) \in M^*/qM^*$.

We show that the oracle $\mathcal{O}$ is used properly, i.e., that what is given to it as input follows a valid distribution $A^{(M)}_{\boldsymbol{s},\psi}$.

Lemma 18. *Let $\varepsilon > 0$ and $\boldsymbol{s} = \Theta(\boldsymbol{x} \bmod qM^*)$. There exists $\psi \in \Psi_{\leq\alpha}$ such that the statistical distance between $A^{(M)}_{\boldsymbol{s},\psi}$ and the distribution of $(\boldsymbol{a}, b)$ is at most 4ε.*

Proof. We first show that the statistical distance between $\boldsymbol{a}$, the first component of each sample, and the uniform distribution on $(R_q)^d$ is at most 2ε. By Lemma 4, the statistical distance between the distribution of $\boldsymbol{z}$ and the uniform distribution on M_q is at most 2ε, because $r \geq q \cdot \eta_\varepsilon(M) = \eta_\varepsilon(qM)$. Then, as Θ^{-1} induces a bijection from M_q to $(R_q)^d$, the statistical distance between the distribution of $\boldsymbol{a} = \Theta^{-1}(\boldsymbol{z} \bmod qM)$ and the uniform distribution on $(R_q)^d$ is at most 2ε.

Now, we show that b is of the shape $b = \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s} \rangle + f$, where f distributed from $D_{\boldsymbol{r}'}$ with $r'_i \leq \alpha$ for all i. We have:

$$b = \frac{1}{q}\langle \boldsymbol{z}, \boldsymbol{y} \rangle + e' = \frac{1}{q}\langle \boldsymbol{z}, \boldsymbol{x} + \boldsymbol{e} \rangle + e' = \frac{1}{q}\langle \boldsymbol{z}, \boldsymbol{x} \rangle + \langle \frac{1}{q}\boldsymbol{z}, \boldsymbol{e} \rangle + e'.$$

By definition, we have $\boldsymbol{z} = \Theta(\boldsymbol{a}) = \sum_{i=1}^{d} (t_i \cdot a_i) \cdot \boldsymbol{b}_i \bmod qM$ with $t_i \in I_i$ and $a_i \in R_q$. By Lemma 1, we have $M^* = \sum_{i=1}^{d} I_i^* \cdot \boldsymbol{b}_i^*$. Then, let $\boldsymbol{x} = \sum_{i=1}^{d} x_i \cdot \mathbf{b}_i^*$, it implies that $x_i \in I_i^* = I_i^{-1} \cdot R^*$ for all i. We also have $\langle \boldsymbol{b}_i, \boldsymbol{b}_j^* \rangle = 1$ if $i = j$ and $\langle \boldsymbol{b}_i, \boldsymbol{b}_j^* \rangle = 0$ otherwise. Then, modulo qR^*:

$$\langle \boldsymbol{z}, \boldsymbol{x} \rangle = \sum_{i,j=1}^{d} (t_i \cdot a_i) \cdot x_j \cdot \langle \boldsymbol{b}_i, \boldsymbol{b}_j^* \rangle = \sum_{i=1}^{d} (t_i \cdot a_i) \cdot x_i = \sum_{i=1}^{d} a_i \cdot (t_i \cdot x_i).$$

Because $\boldsymbol{s} = \Theta(\boldsymbol{x} \bmod qM^*) = (t_1 \cdot x_1 \bmod qR^*, \ldots, t_d \cdot x_d \bmod qR^*)$, we have:

$$\langle \boldsymbol{a}, \boldsymbol{s} \rangle = \sum_{i=1}^{d} a_i \cdot (t_i \cdot x_i) = \langle \boldsymbol{z}, \boldsymbol{x} \rangle \mod qR^*.$$

As a consequence, we obtain that $\frac{1}{q}\langle \boldsymbol{z}, \boldsymbol{x} \rangle = \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s} \rangle \bmod R^*$.

We now show that, conditioned on $\boldsymbol{a}$, the quantity $\langle \frac{1}{q}\boldsymbol{z}/q, \boldsymbol{e} \rangle + e'$ has distribution $\nu_{\boldsymbol{r}'}$ with $r'_i \leq \alpha$ for all i. First, let us analyze the distribution of $\boldsymbol{z'} = \frac{1}{q}\boldsymbol{z}$ knowing $\boldsymbol{a}$. We know that $\boldsymbol{z}$ has distribution $D_{M,r}$ and that $\boldsymbol{a} = \Theta^{-1}(\boldsymbol{z} \bmod qM)$. Let $\boldsymbol{u} = \Theta(\boldsymbol{a}) \bmod qM$, then the residual distribution of $\boldsymbol{z'} = \frac{1}{q}\boldsymbol{z}$ knowing $\boldsymbol{a}$ is $D_{M+\boldsymbol{u}/q, r/q}$ (with $r/q \geq \sqrt{2}\eta_\varepsilon(M)$).

We now show that e' is following the same distribution as $\langle \boldsymbol{e''}, \boldsymbol{e} \rangle$ with $\boldsymbol{e''} \sim \nu_{\boldsymbol{s},\ldots,\boldsymbol{s}}$, $\boldsymbol{s} = (s_i)_i$ and $s_i = s_{i+n/2} = \alpha/\sqrt{2\sum_{k=1}^{d} |\sigma_i(e_k)|^2}$ for $1 \leq i \leq n/2$. By Lemma 10, as the vector $\boldsymbol{e''}$ is distributed from $\nu_{\boldsymbol{s},\ldots,\boldsymbol{s}}$ and $\boldsymbol{e} \in K^d$ is fixed, we have that $\langle \boldsymbol{e''}, \boldsymbol{e} \rangle$ has distribution $\nu_{\boldsymbol{s}'}$ with $s'_i = s'_{i+n/2} = s_i\sqrt{\sum_{k=1}^{d} |\sigma_i(e_k)|^2} = \frac{\alpha}{\sqrt{2}}$.

We are now led to considering the distribution of $\langle \boldsymbol{z'} + \boldsymbol{e''}, \boldsymbol{e} \rangle$. We write $\boldsymbol{e''} = \boldsymbol{e''_1} + \boldsymbol{e''_2}$ with $\boldsymbol{e''_1} \sim \nu_{\alpha/(\sqrt{2}\delta)}$ and $\boldsymbol{e''_2} \sim \nu_{\boldsymbol{s}''}$ with $(s''_i)^2 = s_i^2 - \alpha^2/(2\delta^2)$ (which is positive by assumption on $\|\boldsymbol{e}\|_{2,\infty}$). As we have $\alpha/(\sqrt{2}\delta) = r/q$ and $r/q \geq \sqrt{2}\eta_\varepsilon(M)$, Lemma 6 gives us that the statistical distance between the distribution of $\boldsymbol{z'} + \boldsymbol{e''_1}$ and $\nu_{\alpha/\delta}$ is at most 4ε. As a consequence, the statistical distance between the distribution of $\boldsymbol{z'} + \boldsymbol{e''_1} + \boldsymbol{e''_2}$ and $\nu_{\boldsymbol{r}'',\ldots,\boldsymbol{r}''}$ is at most 4ε, with

$$(r''_i)^2 = \frac{\alpha^2}{\delta^2} + (s''_i)^2 = \frac{\alpha^2}{\delta^2} + s_i^2 - \frac{\alpha^2}{2\delta^2} = \frac{\alpha^2}{2\sum_{k=1}^{d} |\sigma_i(e_k)|^2} + \frac{\alpha^2}{2\delta^2}.$$

By using Lemma 10 again with the fixed vector $\boldsymbol{e}$, we obtain that the statistical distance between the distribution of $\langle \boldsymbol{z} + \boldsymbol{e''}, \boldsymbol{e} \rangle$ and $\nu_{\boldsymbol{r}'}$ is at most 4ε, where

$$r'_i = \sqrt{\frac{\alpha^2}{2} + \frac{\alpha^2 \sum_{k=1}^{d} |\sigma_i(e_k)|^2}{2\delta^2}}.$$

Since $\delta \geq \sqrt{\sum_{k=1}^{d} |\sigma_i(e_k)|^2}$, we have $r'_i \leq \alpha$, as desired. □

As the input of $\mathcal{O}$ is within statistical distance $\leq 4\varepsilon$ from $A^{(M)}_{\boldsymbol{s},\psi}$ for a distribution $\psi \in \Psi_{\leq\alpha}$ and $\boldsymbol{s} = \Theta(\boldsymbol{x} \bmod qM^*)$, oracle $\mathcal{O}$ succeeds with probability at least $p - 4\varepsilon$, where p is its success probability given a valid input distribution. If it does succeed, then the output of our reduction is $\boldsymbol{x} \bmod qM^*$, which completes the proof of Lemma 17. □

A reduction from M-CLWE to M-DLWE. Finally, we describe a reduction from the search version M-CLWE to the decision version M-DLWE. We will use the line of proof of [14] for reducing R-CLWE to R-DLWE. Let $q = 1 \bmod 2n$ be prime. Then $(q) = \prod_{i=1}^{n} \mathfrak{q}_i$ where any $\mathfrak{q}_i$ is a prime ideal with norm $N(\mathfrak{q}_i) = q$. We define the following:

- $\mathfrak{q}_i$**-MLWE**$_{q,\Psi}$, with parameters Ψ a family of distributions over $\mathbb{T}_{R^*}$ and $i \leq n$: Given access to an oracle sampling from $A^{(M)}_{\boldsymbol{s},\psi}$ for some arbitrary $\boldsymbol{s} \in (R_q^*)^d$ and $\psi \in \Psi$, find $\boldsymbol{s} \bmod \mathfrak{q}_i R_q^*$.
- **Hybrid distribution** $A^{(M,i)}_{\boldsymbol{s},\psi}$, with parameters ψ a distribution over $\mathbb{T}_{R^*}$, $\boldsymbol{s} \in (R_q^*)^d$ and $i \leq n$: The distribution $A^{(M,i)}_{\boldsymbol{s},\psi}$ over $(R_q)^d \times \mathbb{T}_{R^*}$ is defined as follows: Choose $(\boldsymbol{a}, b)$ from $A^{(M)}_{\boldsymbol{s},\psi}$ and return $(\boldsymbol{a}, b + r/q)$ where $r \in R_q^*$ is uniformly random and independent in $R_q^*/\mathfrak{q}_j R^*$ for all $j \leq i$, and is 0 modulo the remaining $\mathfrak{q}_j R^*$'s.

- **DecMLWE**$^i_{q,\Psi}$, with parameters Ψ a family of distributions on $\mathbb{T}_{R^*}$ and $i \leq n$: Given access to an oracle sampling from $A^{(M,j)}_{\boldsymbol{s},\psi}$ for arbitrary $\boldsymbol{s} \in (R_q^*)^d$, $\psi \in \Psi$ and $j \in \{i-, i\}$, find j.

We consider the following sequence of reductions:

$$\text{MLWE}_{q,\Psi} \xrightarrow{\text{Lemma 19}} \mathfrak{q}_i\text{-MLWE}_{q,\Psi} \xrightarrow{\text{Lemma 20}} \text{DecMLWE}^i_{q,\Psi} \xrightarrow{\text{[14, Le. 5.11 \& 5.13]}} \text{DMLWE}_{q,\Upsilon}$$

We only explicit the first two reductions, as the last one carries over directly from the ring setting [14, Le. 5.11 & 5.13] to the module setting (the proof randomizes the noise distribution Ψ, which is the same in the ring an module settings). To show the two new lemmata, we adapt proofs from [14].

Lemma 19. *For any $i \in [n]$ and any family Ψ closed under all the automorphisms of K, there exists a probabilistic polynomial time reduction from* M-CLWE$_{q,\Psi}$ *to* $\mathfrak{q}_i$-MLWE$_{q,\Psi}$*.*

Proof. We aim at using an oracle solving $\mathfrak{q}_i$-MLWE for finding the values of $\boldsymbol{s}$ mod $\mathfrak{q}_j R^*$ for every $j \in [n]$. Then, by the Chinese Reminder Theorem, this allows us to construct $\boldsymbol{s}$ mod R^* and to solve M-CLWE.

We use the K-automorphisms, defined by $\tau_k(\xi) = \xi^k$ for all $k \in [n]$. We choose $k \in [n]$ such that $t_k(\mathfrak{q}_j) = \mathfrak{q}_i$. The reduction is as follows:

- For every sample $(\boldsymbol{a}, b) \leftarrow A^{(M)}_{\boldsymbol{s},\psi}$, create the sample $(\boldsymbol{a}', b')$ with $\boldsymbol{a}' = (\tau_k(a_1), \ldots, \tau_k(a_d))$ and $b' = \tau_k(b)$.
- Use the oracle of $\mathfrak{q}_i$-MLWE with these samples, and get $\boldsymbol{t} \in (R^*/\mathfrak{q}_i R^*)^d$.
- Return $(\tau_k^{-1}(t_1), \ldots, \tau_k^{-1}(t_d)) \in (R^*/\mathfrak{q}_j R^*)^d$.

We show that $\tau_k^{-1}(t_l) = s_l$ mod $\mathfrak{q}_j R^*$ for all $l \in [d]$. By definition, we have $b = \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s} \rangle + e$ mod R^* with $\langle \boldsymbol{a}, \boldsymbol{s} \rangle = \sum_{i=1}^d a_i \cdot s_i$. As a consequence, we have:

$$b' = \tau_k(b) = \frac{1}{q}\sum_{i=1}^{d} \tau_k(a_i) \cdot \tau_k(s_i) + \tau_k(e) = \frac{1}{q}\langle \boldsymbol{a}', \boldsymbol{s}' \rangle + \tau_k(e) \bmod R^*,$$

with $\boldsymbol{s}' = (\tau_k(s_1), \ldots, \tau_k(s_d))$. As τ_k is an automorphism, the vector $\boldsymbol{a}'$ is uniformly distributed in $(R_q)^d$. Also, as Ψ is closed under the automorphisms of K, we have $\psi' := \tau_k(\psi) \in \Psi$. Overall, the pairs $(\boldsymbol{a}', b')$ are distributed as $A^{(M)}_{\boldsymbol{s}',\psi'}$. If successful, the $\mathfrak{q}_i$-MLWE oracle outputs $\boldsymbol{t} = \boldsymbol{s}'$ mod $\mathfrak{q}_i R^* = (\tau_k(s_1)$ mod $\mathfrak{q}_i R^*, \ldots, \tau_k(s_d)$ mod $\mathfrak{q}_i R^*)$. Then our reduction returns $(\tau_k^{-1}(t_1), \ldots, \tau_k^{-1}(t_d)) \in (R^*/\mathfrak{q}_j R^*)^d$, which is equal to $\boldsymbol{s}$ mod $\mathfrak{q}_j R^*$. □

Lemma 20. *For any $i \in [n]$, there exists a probabilistic polynomial time reduction from $\mathfrak{q}_i$-MLWE$_{q,\Psi}$ to* DecMLWE$^i_{q,\Psi}$*.*

Proof. We want to find $\boldsymbol{s}$ mod $\mathfrak{q}_i R^*$ from samples from $A^{(M)}_{\boldsymbol{s},\psi}$, by using an oracle that solves the DecMLWE$^i_{q,\Psi}$ problem. The principle of the proof is to find, one by one, each one of the d coordinates of $\boldsymbol{s}$ mod $\mathfrak{q}_i R^*$ by using the oracle of DecMLWE$^i_{q,\Psi}$. For each coordinate, there are $N(\mathfrak{q}_i) = q \leq \text{poly}(n)$ possibilities. Therefore, it is possible to try them all in order to find the correct one. To check that a guess is correct, we use the same approach as in [31, Le. 4.2] and randomize a coordinate of $\boldsymbol{a}$.

To find s_1 mod $\mathfrak{q}_i R^*$, we proceed as follows. Let $(\boldsymbol{a}, b)$ be distributed as $A^{(M)}_{\boldsymbol{s},\psi}$ and let $x \in R_q^*$; we want to know if $x = s_1$ mod $\mathfrak{q}_i R^*$. We construct the following pair:

$$(\boldsymbol{a}', b') := \left(\boldsymbol{a} + (y, 0, \ldots, 0), b + \frac{1}{q}(r + xy)\right),$$

where $y \in R_q$ is sampled uniformly modulo $\mathfrak{q}_i$, and is 0 modulo all the remaining $\mathfrak{q}_j$'s, and where $r \in R_q^*$ is uniformly random and independent modulo $\mathfrak{q}_j R^*$ for all $j < i$, and 0 modulo all the remaining $\mathfrak{q}_j R^*$'s.

Now, we show that if $x = s_1 \bmod \mathfrak{q}_i R^*$, then the pair $(\boldsymbol{a}', b')$ is distributed from $A^{(M,i-)}_{\boldsymbol{s},\psi}$ and if $x \neq s_1 \bmod \mathfrak{q}_i R^*$, it is distributed from $A^{(M,i)}_{\boldsymbol{s},\psi}$. First, notice that the vector $\boldsymbol{a}'$ is uniformly distributed in $(R_q)^d$. Now, we write b' as follows:

$$b' = b + \frac{1}{q}(r + xy) = \frac{1}{q}\left(\sum_{i=1}^{d} a_i \cdot s_i + r + xy\right) + e = \left(\frac{1}{q}\langle \boldsymbol{a}', \boldsymbol{s}\rangle + e\right) + \frac{1}{q}\left(r + y(x - s_1)\right).$$

We have two cases:

- If $x = s_1 \bmod \mathfrak{q}_i R^*$, then by the Chinese Reminder Theorem we have $y(x - s_1) = 0 \in R_q^*$. As r is chosen uniformly random and independent modulo $\mathfrak{q}_j R^*$ for all $j < i$, and is 0 modulo all the remaining $\mathfrak{q}_j R^*$'s, we obtain that the pair $(\boldsymbol{a}', b')$ has distribution $A^{(M,i-)}_{\boldsymbol{s},\psi}$.
- If $x \neq s_1 \bmod \mathfrak{q}_i R^*$, then $y(x - s_1)$ is uniformly distributed modulo $\mathfrak{q}_i R^*$, because $R^*/\mathfrak{q}_i R^*$ is a field (the ideal $\mathfrak{q}_i$ maximal). Also, the quantity $y(x - s_1)$ is 0 modulo the other $\mathfrak{q}_j R^*$'s. As a consequence, we have that $(r + y(x - s_1))$ is uniformly random and independent modulo $\mathfrak{q}_j R^*$ for all $j \leq i$ and is 0 modulo all the remaining $\mathfrak{q}_j R^*$'s. We obtain that the pair $(\boldsymbol{a}', b')$ is distributed as $A^{(M,i)}_{\boldsymbol{s},\psi}$.

We repeat this process d times (once for each coordinate of $\boldsymbol{s}$), to obtain $\boldsymbol{s} \bmod \mathfrak{q}_i R^*$. □

This completes the proof of the Theorem 8.

5 Converse reductions

In this section, we show that M-SIS and M-LWE both reduce to Mod-GIVP. This provides converse results to Theorems 4 and 8.

Reducing M-SIS *to* Mod-GIVP. In this paragraph, we identify elements of the field K with their polynomial representations. Let $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_m$ be sampled uniformly and independently in R_q^d. Finding $\boldsymbol{z} = (z_1 | \ldots | z_m) \in R^m \setminus \boldsymbol{0}$ such that $\sum_i z_i \boldsymbol{a}_i = 0 \bmod q$ and $\|\boldsymbol{z}\| \leq \beta$ corresponds to finding a short vector in the lattice:

$$\boldsymbol{A}^{\perp} = \left\{\boldsymbol{y} \in R^m : \boldsymbol{A}^T \boldsymbol{y} = 0 \bmod q\right\},$$

where $\boldsymbol{A} \in R_q^{d \times m}$ is the matrix whose rows are the $\boldsymbol{a}_i$'s. As this lattice is a module lattice, if we solve Mod-GIVP$^{\eta_\varepsilon}_{\gamma}$ given as input an arbitrary basis of $\boldsymbol{A}^{\perp}$ (which can be computed efficiently given $\boldsymbol{A}$), then we obtain a solution to the M-SIS instance, for $\beta = \gamma \cdot \eta_\varepsilon(\boldsymbol{A}^{\perp})$. To assess the effectiveness of this reduction from M-SIS to Mod-GIVP, we are thus led to estimating $\eta_\varepsilon(\boldsymbol{A}^{\perp})$ for $\boldsymbol{A}$ sampled uniformly in $R_q^{m \times d}$. For this task, it is classical to study the dual lattice, as we have $\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2k(1+1/\varepsilon))}{\pi}}/\lambda_1^\infty(\Lambda^*)$ for any k-dimensional lattice Λ (see Lemma 3). By [32, Le. 7], the dual of the lattice $\boldsymbol{A}^{\perp}$ is $\frac{1}{q} L_q(\boldsymbol{A}^{\times})$ where $\boldsymbol{A}^{\times}$ is the matrix obtained by replacing each entry $a_{i,j}(x) \in R_q$ of $\boldsymbol{A}$ by $a_{i,j}(x^{-1})$, and

$$L_q(\boldsymbol{B}) = \left\{\boldsymbol{y} \in R^m : \exists \boldsymbol{s} \in R_q^d, \boldsymbol{B}\boldsymbol{s} = \boldsymbol{y} \bmod q\right\}, \quad \text{for any} \quad \boldsymbol{B} \in R_q^{m \times d}.$$

Note that the map $a(x) \mapsto a(x^{-1})$ is an isomorphism. As a consequence, it suffices to obtain a probabilistic lower bound on $\lambda_1^\infty(L_q(\boldsymbol{A}))$, for $\boldsymbol{A}$ uniform in $R_q^{m \times d}$.

Similarly, for reducing M-SIS to M-SIVP, one is led to bounding $\lambda_{mn}(\boldsymbol{A})^{\perp}$. As $\lambda_k(\Lambda) \leq k/\lambda_1(\Lambda) \leq k^{3/2}/\lambda_1^\infty(\Lambda^*)$ for any k-dimensional Λ, it is also sufficient to obtain an lower bound for $\lambda_1^\infty(\Lambda^*)$.

Lemma 21. *Let n, m, d, q be positive integers with $d \leq m$. We have:*

$$\Pr_{A \hookleftarrow U(R_q^{m \times d})}\left[\lambda_1^\infty(L_q(\boldsymbol{A})) \geq \frac{1}{8\sqrt{n}} q^{1 - \frac{d}{m}}\right] \geq 1 - \left(\frac{1}{2\sqrt{n}}\right)^{nd}.$$

Proof. We obtain this result by generalizing the proof of [32, Le. 8]. By the union bound, the probability that $L_q(\boldsymbol{A})$ contains a non-zero vector of infinity norm $\leq B := \frac{1}{8\sqrt{n}} q^{1-\frac{d}{m}}$ is bounded from above by:

$$\sum_{\substack{\boldsymbol{t} \in R_q^m \\ 0 < \|\boldsymbol{t}\|_\infty \leq B}} \sum_{\boldsymbol{s} \in R_q^d} \Pr_{\boldsymbol{A} \hookleftarrow U(R_q^{m \times d})}[\boldsymbol{A}\boldsymbol{s} = \boldsymbol{t}] \;=\; \sum_{\substack{\boldsymbol{t} \in R_q^m \\ 0 < \|\boldsymbol{t}\|_\infty \leq B}} \sum_{\boldsymbol{s} \in R_q^d} \prod_{i \leq m} \Pr_{\boldsymbol{a} \hookleftarrow U(R_q^d)}[\langle \boldsymbol{a}, \boldsymbol{s} \rangle = t_i].$$

We now consider the probability (over the randomness of $\boldsymbol{a}$) that $\langle \boldsymbol{a}, \boldsymbol{s} \rangle = t_i$. For this purpose, we consider the decomposition of R_q as a Cartesian product of finite fields. If $x^n + 1 = \prod_k \Phi_k \bmod q$ with irreducible polynomials Φ_k, then these Φ_k's have a common degree δ dividing n. Then we have $R_q \simeq (\mathbb{F}_{q^\delta})^{\frac{n}{\delta}}$, where $\mathbb{F}_{q^\delta}$ is the field with q^δ elements. This ring isomorphism can be explicited: It is given by the Chinese Remainder Theorem map $x \mapsto (x \bmod \Phi_1, \ldots, x \bmod \Phi_{n/\delta})$. Now, the equality $\langle \boldsymbol{a}, \boldsymbol{s} \rangle = t_i$ holds if and only if it holds over all CRT components. Wlog we consider Φ_1. If t_i and all the coordinates of $\boldsymbol{s}$ are zero modulo Φ_1, then the probability that $\langle \boldsymbol{a}, \boldsymbol{s} \rangle = t_i \bmod \Phi_1$ is 1. Otherwise, if t_i or some coordinate of $\boldsymbol{s}$ is non-zero on that component, then the probability is $\leq q^\delta$. As a consequence, the probability under scope is bounded from above by:

$$\sum_{0 \leq k \leq n/\delta} \sum_{\substack{h = \prod_{i \in S'} \Phi_i \\ S' \subseteq S \\ |S'| = k}} \sum_{\substack{\boldsymbol{s} \in R_q^d \\ \forall i, h | s_i}} \sum_{\substack{\boldsymbol{t} \in R_q^m \\ 0 < \|\boldsymbol{t}\|_\infty \leq B \\ \forall i, h | t_i}} q^{m(k\delta - n)} \;\leq\; \sum_{0 \leq k \leq n/\delta} \sum_{\substack{h = \prod_{i \in S'} \Phi_i \\ S' \subseteq S \\ |S'| = k}} \sum_{\substack{\boldsymbol{t} \in R_q^m \\ \|\boldsymbol{t}\|_\infty \leq B \\ \forall i, h | t_i}} q^{(m-d)(k\delta - n)}.$$

The rest of the proof is as in [32]. □

As a consequence of the result above and the preceding discussion, we obtain the following converse to Theorem 4. Note that even for $d = 1$ (i.e., for an R-SIS instance), the resulting Mod-GIVP instance has module rank m: This result does not provide a reduction from R-SIS to Id-GIVP.

Theorem 9. *For any $d \geq 1$ and $\varepsilon(nd) = (nd)^{-\omega(1)}$, there is a probabilistic polynomial time reduction from solving* M-SIS$_{q,m,\beta}$ *to solving* Mod-GIVP$^{\eta_\varepsilon}_{\gamma}$ *(with module rank m), for any $m(nd), q(nd)$, $\beta(dn)$ and $\gamma(nd)$ such that $\beta \geq \gamma\sqrt{n}\omega\left(\sqrt{\log(1/\varepsilon)}\right) \cdot q^{\frac{d}{m}}$ and $m, \log q \leq \mathrm{poly}(nd)$.*

Reducing M-DLWE *to* Mod-GIVP. One of the classical ways for solving LWE consists in solving an associated SIS instance [22]. We propose an adaptation of this approach to module lattices: We reduce M-DLWE to M-SIS and then combine this reduction with Theorem 9 above.

Let us sample $\boldsymbol{s}$ uniformly in R_q^d, and ψ from Υ_α. More precisely, we sample x_i from $\Gamma(2,1)$ for $i \leq n/2$, define $r_i = r_{i+n/2} = \alpha\sqrt{1 + \sqrt{n} x_i}$, and let $\psi = \nu_{\boldsymbol{r}}$. Assume that we have access to arbitrarily many samples $(\boldsymbol{a}_i, b_i) \in R_q^d \times \mathbb{T}_{R^*}$ with a_i uniform in R_q^d and all the b_i's uniform and independent in $\mathbb{T}_{R^*}$, or all the b_i's of the form $b_i = \frac{1}{q}\langle \boldsymbol{a}_i, \boldsymbol{s} \rangle + e_i$ with the e_i's sampled from ψ. Our goal is to determine with noticeable advantage which situation we are in.

We consider m such samples (with m to be optimized later). Let $\boldsymbol{A} \in R_q^{m \times d}$ be the matrix whose rows are the $\boldsymbol{a}_i$'s. By solving M-SIS$_{q,m,\beta}$ for $\boldsymbol{A}^t$, we obtain a non-zero vector $\boldsymbol{z} \in R^m$ such that $\|\sigma_P(\boldsymbol{z})\| \leq \beta$ and $\boldsymbol{z}^t \cdot \boldsymbol{A} = 0 \bmod q$. Now, we compute $\langle \boldsymbol{z}, \boldsymbol{b} \rangle$, where $\boldsymbol{b} \in \mathbb{T}^m_{R*}$ is the vector made of the b_i's. If the b_i's are uniform independent of the $\boldsymbol{a}_i$'s, then the inner product $ps\boldsymbol{z}\boldsymbol{b}$ is uniformly distributed in $\mathbb{T}_{R^*}$. Otherwise, we have $\langle \boldsymbol{z}, \boldsymbol{b} \rangle = \langle \boldsymbol{z}, \boldsymbol{e} \rangle$ (modulo R^*), where $\boldsymbol{e}$ is the vector made of the e_i's. By Lemma 10, we have that $\langle \boldsymbol{z}, \boldsymbol{e} \rangle$ is distributed as $\nu_{\boldsymbol{r}'}$ with $r'_j = r_j \cdot \sqrt{\sum_{k \leq m} |\sigma_j(z_k)|^2}$ for all $j \leq n$. As a consequence, we have

$$\begin{aligned} \|\sigma_C(\langle \boldsymbol{z}, \boldsymbol{b} \rangle)\| \;&=\; \|\sigma_H(\langle \boldsymbol{z}, \boldsymbol{b} \rangle)\| \leq t\sqrt{n} \cdot \max_j |r'_j| \\ &\leq t\sqrt{n} \cdot \sigma_C(\boldsymbol{z}) \cdot \max_j |r_j| \leq 2tn^{3/2}\alpha\beta \cdot \max_j |x_j|, \end{aligned}$$

with probability $\geq 1 - 2^{-\Omega(nt^2)}$ over the randomness of the e_i's. Furthermore, as we have than $x_j \leq t$ with probability $\geq 1 - (2+t)e^{-t}$ for all j, we obtain that the bound above is itself smaller than $2t^2 n^{3/2}\alpha\beta$ with probability $\geq 1 - nt2^{-\Omega(t)}$. As $R^* = \frac{1}{n}R$, if the latter upper bound is smaller than $\frac{1}{4n}$, then $\langle \boldsymbol{z}, \boldsymbol{b} \rangle$ will be unexpectedly small.

Overall, we have proved that if β is such that $n^{5/2}\omega(\log(nd)) \cdot \alpha\beta < 1$, then we can distinguish between the two challenge distributions with non-negligible advantage. By Theorem 9, we thus obtain a reduction from Mod-GIVP$_\gamma^{\eta_\varepsilon}$ with module rank m to M-DLWE$_{q,\Upsilon_\alpha}$, if γ is such that $\alpha\gamma n^3 \omega \left(\log(nd)\sqrt{\log(1/\varepsilon)}\right) q^{\frac{d}{m}} < 1$. Taking $m = d \log q$ leads to the following result.

Theorem 10. *For any $d \geq 1$ and $\varepsilon(nd) = (nd)^{-\omega(1)}$, there is a probabilistic polynomial time reduction from solving* M-DLWE$_{q,\Upsilon_\alpha}$ *to solving* Mod-GIVP$_\gamma^{\eta_\varepsilon}$ *(with module rank $d \log q$), for any $\alpha(dn)$ and $\gamma(nd)$ such that $\frac{1}{\alpha} \geq \gamma n^3 \omega \left(\log(nd)\sqrt{\log(1/\varepsilon)}\right)$ and $\log q \leq \mathrm{poly}(nd)$.*

Acknowledgements

We thank Guillaume Hanrot, Oded Regev and Ron Steinfeld for helpful discussions. Significant parts of the writing of this article were undergone while the authors were visiting Macquarie University, whose hospitality is gratefully acknowledged. The authors were partly supported by the Australian Research Council under Discovery Grant DP0987734.

References

1. M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proc. of STOC*, pages 99–108. ACM, 1996.
2. B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Proc. of CRYPTO*, volume 5677 of *LNCS*, pages 595–618. Springer, 2009.
3. J. Blömer and J.-P. Seifert. On the complexity of computing short linearly independent vectors and short bases in a lattice. In *Proc. of STOC*. ACM, 1999.
4. D. Boneh and D. M. Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In *Proc. of PKC*, volume 6571 of *LNCS*, pages 1–16. Springer, 2011.
5. W. Bosma and M. Pohst. Computations with finitely generated modules over Dedekind rings. In *Proc. of ISSAC*, pages 151–156, 1991.
6. Z. Brakerski, C. Gentry, and V. Vaikuntanathan. Fully homomorphic encryption without bootstrapping. *Cryptology ePrint Archive, Report 2011/277*, 2011.
7. H. Cohen. *Advanced topics in computational number theory.* Springer, 2000.
8. C. Fieker and D. Stehlé. Short bases of lattices over number fields. In *Proc. of ANTS-IX*, volume 6197 of *LNCS*, pages 157–173. Springer, 2010.
9. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008.
10. S. D. Gordon, J. Katz, and V. Vaikuntanathan. A group signature scheme from lattice assumptions. In *Proc. of ASIACRYPT*, volume 6477 of *LNCS*, pages 395–412. Springer, 2010.
11. J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In *Proc. of ANTS-III*, pages 267–288, 1998.
12. E. Kiltz, K. Pietrzak, D. Cash, A. Jain, and D. Venturi. Efficient authentication from hard learning problems. In *Proc. of EUROCRYPT*, volume 6632 of *LNCS*. Springer, 2011.
13. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proc. of ICALP (2)*, volume 4052 of *LNCS*, pages 144–155. Springer, 2006.
14. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *Proc. of EUROCRYPT*, pages 1–23, 2010. All theorem numberings used correspond to the draft of the full version, available from the authors upon request.
15. D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *Proc. of FOCS*, pages 356–365. IEEE Computer Society Press, 2002. Conference version of [17].

16. D. Micciancio. Almost perfect lattices, the covering radius problem, and applications to Ajtai's connection factor. *SIAM Journal on Computing*, 34(1):118–169, 2004. Preliminary version in STOC 2002.
17. D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complexity*, 16(4):365–411, 2007. Journal version of [15].
18. D. Micciancio and S. Goldwasser. *Complexity of lattice problems: a cryptographic perspective*. Kluwer Academic Press, 2002.
19. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller, 2011. IACR Cryptology ePrint Archive, report 2011/501.
20. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measure. In *Proc. of FOCS*, pages 371–381. IEEE, 2004. Conference version of [21].
21. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput*, 37(1):267–302, 2007. Journal version of [20].
22. D. Micciancio and O. Regev. Lattice-based cryptography. In D.J. Bernstein, J. Buchmann, and E. Dahmen, editors, *Post Quantum Cryptography*, pages 147–191. Springer, 2009.
23. R. A. Mollin. *Algebraic Number Theory*. Chapman and Hall/CRC Press, 1999.
24. C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Comput. Complexity*, 2(17):300–351, 2008.
25. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.
26. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Proc. of TCC*, volume 3876 of *LNCS*, pages 145–166. Springer, 2006.
27. C. Peikert and A. Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *Proc. of STOC*, pages 478–487. ACM, 2007.
28. K. Pietrzak. Subspace LWE, 2011. Available at `http://homepages.cwi.nl/~pietrzak/publications/SLWE.pdf`.
29. O. Regev. Lecture notes of *lattices in computer science*, taught at the Computer Science Tel Aviv University. Available at `http://www.cs.tau.il/~odedr/`.
30. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proc. of STOC*, pages 84–93, 2005.
31. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009. Journal version of [30].
32. D. Stehlé and R. Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *Proc. of EUROCRYPT*, volume 6632 of *LNCS*, pages 27–47. Springer, 2011. All theorem numberings used correspond to the draft of the full version, available from the authors upon request.
33. D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *Proc. of ASIACRYPT*, volume 5912 of *LNCS*, pages 617–635. Springer, 2009.