

Worst-Case to Average-Case Reductions for Module Lattices

Adeline Langlois and Damien Stehlé

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 bridge 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 $\mathbf{s}_1, \dots, \mathbf{s}_n$ in a given n -dimensional lattice, that minimize $\max_i \|\mathbf{s}_i\|$. (The dimension of a lattice is the dimension of the vector space spanned by its vectors.) A standard relaxation of this optimization problem, parametrized by $\gamma(n) \geq 1$, consists in requesting that $\max_i \|\mathbf{s}_i\|$ 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 [4]). 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 [25].

Lattice-based cryptography is a branch of cryptography exploiting the presumed 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 (the worst-case problem) which is presumed difficult. Note here that a worst-case problem needs every instance to be solved (e.g., with non-negligible probability), as an average-case problem only requires some instances (a non-negligible proportion) to be solved. We refer to [25] 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 nonzero solution $\mathbf{z} \in \mathbb{Z}^m$ to the homogeneous linear system $\mathbf{A}\mathbf{z} = \mathbf{0} \pmod{q}$ for uniformly random $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ (the notation \mathbb{Z}_q denotes the ring of integers modulo q). The second one, introduced by Regev in 2005 [35], 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 \mathbf{s} , given arbitrarily many independent pairs $(\mathbf{a}, \frac{1}{q}\langle \mathbf{a}, \mathbf{s} \rangle + e)$ for a vector $\mathbf{a} \in \mathbb{Z}_q^n$ 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 $(\mathbf{a}_i, \frac{1}{q}\langle \mathbf{a}_i, \mathbf{s} \rangle + e_i)$ are considered, the goal is to find \mathbf{s} from $(\mathbf{A}, \frac{1}{q}\mathbf{A}\mathbf{s} + \mathbf{e})$, where the rows of \mathbf{A} correspond to the \mathbf{a}_i 's and $\mathbf{e} = (e_1, \dots, e_m)^T$. The decision counterpart of LWE consists in distinguishing between arbitrarily many independent pairs $(\mathbf{a}, \frac{1}{q}\langle \mathbf{a}, \mathbf{s} \rangle + e)$ 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 [35,36] 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 \mathbf{A} .

To circumvent this inherent inefficiency, Micciancio [18,20] — inspired from the efficient NTRU encryption scheme [13] that can itself 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 \mathbf{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 [20] was later replaced by a more powerful variant [15,31], 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. [16] (see also [39]). Similarly to SIS and LWE, these problems admit reductions from worst-case lattice problems [15,31,16], 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_γ .

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 [7], 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 in the proof of Theorem 3.6, and a (quantum) reduction from Mod-SIVP to M-LWE in both its search and decision versions in the proofs of Theorems 4.7 and 4.8.

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 ([11] and [15] for M-SIS, and [36] and [16] 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.

On the way, we improve the state-of-the-art results on the hardness of R-SIS and R-LWE. Concerning R-SIS: We improve the reduction from Id-SIVP_γ by allowing for smaller values of q ; this improvement is obtained by adapting a technique based on the Chinese Remainder Theorem and developed by Lyubashevsky et al. in [16] in the context of R-LWE; its application to R-SIS was suggested in [16] but left open. Concerning R-LWE: We show that R-LWE is hard for all sufficiently large q , independently of the arithmetic properties of q with respect to the ring dimension n ; this improvement is obtained by adapting the modulus-switching technique developed by Brakerski et al. in [8] in the context of LWE.

A larger toolbox for cryptographic design. The hardness results for M-SIS and M-LWE possibly enlarge the toolbox 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:

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

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:

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

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 no easier than 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 no easier than 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. d th power) of the lower bound provided by relying on R-SIS. This assertion relies on the reasonable conjecture that the module structure is harder to exploit in SIVP solvers than the ideal structure (because an ideal lattice can be embedded in a module lattice of any dimension).

From the cryptographic construction viewpoint, we expect that most constructions based on R-SIS and R-LWE can be adapted to M-SIS and M-LWE, with an efficiency slowdown (in terms of memory requirements, communication costs and algorithm run-times) bounded by a constant factor when $d = O(1)$.

Hedging against a possible non-hardness of Id-SIVP. 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 security relies on Mod-SIVP instead of Id-SIVP. We argue here that Mod-SIVP is possibly a harder problem than Id-SIVP.

First, 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, hinting that Id-SIVP is possibly easier to solve than R-SIS/R-LWE. We show (in Section 5) that the Mod-SIVP to M-SIS/M-LWE reductions admit converse reductions. Further, M-LWE and M-SIS are obviously no easier than R-LWE and R-SIS. Similarly, Mod-SIVP can trivially be shown to be no easier than Id-SIVP (as any Id-SIVP instance can be embedded into a Mod-SIVP instance of higher dimension). As a result, the existing reductions hint that Mod-SIVP, M-LWE and M-SIS are possibly harder than R-LWE and R-SIS, which are themselves possibly harder than Id-SIVP.

Second, 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, and fail to hold as soon as the rank d of the module is greater than 1 (i.e., a phase transition between $d = 1$ and $d > 1$). Such weaknesses due to the multiplicative structure actually exist for some lattice problems. Consider for example the task of estimating, within a factor γ , the euclidean norm of the shortest nonzero vector in the lattice (known as GapSVP $_\gamma$). This problem is suspected to be extremely hard in the worst case for values of γ that are polynomial in the lattice dimension. But it is easy for ideal lattices, as Minkowski's bound on the lattice minimum is known to be essentially sharp in that case (see, e.g., [32, Se. 6]).

Further, we suspect that this problem is hard in the worst case for module lattices with module rank greater than 1, as it would allow one to efficiently solve R-LWE. We briefly describe that reduction. With the same notations as above, R-LWE with two sample pairs consists in deciding whether a_1, b_1, a_2, b_2 are uniformly and independently sampled in $\mathbb{Z}_q[x]/(x^n + 1) \times \mathbb{Z}_q[x]/(x^n + 1)$, or whether there exists $s \in \mathbb{Z}_q[x]/(x^n + 1)$ such that both $b_1 - a_1 \cdot s$ and $b_2 - a_2 s$ have small coefficients. In the first case, the shortest nonzero vectors of the module lattice $(a_1, a_2) \cdot \mathbb{Z}_q[x]/(x^n + 1) + (b_1, b_2) \cdot \mathbb{Z}_q[x]/(x^n + 1) + (q\mathbb{Z}[x]/(x^n + 1))^2$ are expected to be of norms that are close to Minkowski's bound (i.e., around $\sqrt{n}q^{1/2}$, up to a constant factor). In the second case, there is an unexpectedly short nonzero vector in it. We note that no such phase transition is known for Id-SIVP, but it cannot be ruled out given our current knowledge.

Related works. Most SIVP to SIS reductions (including ours) consider the euclidean norm. Peikert [28] 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 \mathbf{s} from the same distribution as the errors [3], in sampling the error vectors from other distributions [29,12] and in relaxing the conditions on the factorisation of the modulus [22, Se. 3] (see also the references therein). Other cryptographically useful variants of SIS and LWE proven as secure as SIVP include k -SIS [5], ISIS [11], subspace-LWE [14,33] and extended-LWE [27,2,8].

In [29,8], Peikert and Brakerski et al. partially dequantized Regev's proof of hardness of LWE [36], by proposing a reduction from the decisional GapSVP $_\gamma$ problem to LWE. 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 dequantization carries over to the module case, by giving a reduction from GapSVP restricted to module lattices to M-LWE (using Lemma 4.14 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). The reduction of Brakerski et al. consists of several steps, the first one being Peikert's reduction. It is thus equally useless in the case of ideal lattices.

Some computational aspects of module lattices have been investigated in [6,10] (see also [9, 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.

Peikert and Rosen [32] observed that solving R-SIS exactly consists in finding a short nonzero vector in a module lattice.

Open problems. The assumption that standard lattice problems remain hard for ideal lattices is extensively used in modern cryptographic constructions, via the R-SIS and R-LWE problems. However, ideal lattices are very structured objects, and the existence of algorithms exploiting their specific properties should not be discarded too quickly. The fact that GapSVP is easy for these lattices should be seen as a warning of potential risk. The hardness of lattice problems restricted on ideal lattices remains to be assessed. Another approach would be to base the hardness of R-SIS and R-LWE on the worst-case hardness of standard lattice problems for more general classes of lattices. Note that a result in that direction is obtained in [8], but only in the case of an exponential modulus q , which is of limited practical interest.

No weakness is currently known for module lattices of module rank greater than 1. Designing algorithms that exploit this structure is a very exciting research problem, as it may also impact the hardness of R-SIS and R-LWE, as explained above.

Finally, it could be of interest to adapt the main result of [8] to the M-LWE/R-LWE setting. Brakerski et al. present a self-reduction for LWE that trades modulus bitsize for dimension: LWE with modulus q and dimension n is (essentially) as hard as LWE with modulus $\approx q^k$ and dimension n/k .

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

We use standard Landau notations. Furthermore, we say that a function $f(n)$ is $\text{poly}(n)$ 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 $2^{-\Omega(n)}$.

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 any finite measure 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 n for any function $\gamma(n) \in \text{poly}(n)$. This gap between the costs of the algorithms involved in the cryptographic primitives and the cost of the best known attacks is a very attractive feature. For this reason, it is interesting to consider two regimes when analyzing worst-case to average-case reductions from SIVP: In the first regime, we assume we are given a polynomial time algorithm for solving the considered average-case problem with non-negligible probability; In the second (somewhat less standard) regime, we assume we are given a subexponential-time algorithm for solving the considered average-case problem with non-exponentially small probability. The precise parameters of the problems at both ends of the reductions sometimes have to be set differently, so that the reductions essentially preserve the success probabilities and distinguishing advantages. We concentrate here on the first regime, but give the parameter adjustments for the main results to hold in the second setting.

Polynomial representation versus canonical embedding. We propose a unified analysis of R-SIS/M-SIS and R-LWE/M-LWE by only considering the complex canonical embeddings of the ring elements. Note that all prior works on R-SIS except [32] used the polynomial embedding. However, the canonical embedding representation is mathematically sounder, and the unification leads to a more natural connection between R-SIS and R-LWE.

2.1 Some algebraic number theory

In the following, we recall a few facts on elementary algebraic theory in this special case. We refer the reader to [26] and [16,17] for thorough introductions to the topics covered in this subsection.

Number field and cyclotomic field. 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 belongs to $\mathbb{Z}[X]$. 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 $\{1, \xi, \dots, \xi^{n-1}\}$, where n is the degree of f . We call n the degree of K . 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 $\mathbb{Z}[\xi] = \sum_{j=1}^n \mathbb{Z} \cdot \xi^j \subseteq R$. In general, this inclusion can be strict.

A cyclotomic field is a field $K = \mathbb{Q}(\xi)$ where ξ is a root of unity. If ξ is a primitive ν -th root of unity, then it is a root of the ν -th cyclotomic polynomial Φ_ν . The degree $n = \phi(\nu)$ of Φ_ν is the degree of K (here $\phi(\cdot)$ denotes Euler's totient function). In the case of cyclotomic fields, we have $R = \mathbb{Z}[\xi]$.

In this work, all number fields will be cyclotomic fields.

Complex embeddings. The canonical embeddings are the n ring homomorphisms $\sigma_j : K \rightarrow \mathbb{C}$ that fix every element of \mathbb{Q} . In our particular case of cyclotomic fields, all n embeddings are complex: They are defined by $\sigma_j : \xi \mapsto \xi^j$ for any $j \in \mathbb{Z}_\nu^\times$. Note that if j is invertible modulo ν , then so is $\nu - j$, and $\sigma_{\nu-j} = \overline{\sigma_j}$. For notational simplicity, we let \mathbb{J} denote $[\nu/2] \cap \mathbb{Z}_\nu^\times$. We call *canonical embedding vector* the ring homomorphism $\sigma_C : K \rightarrow \mathbb{C}^n$ defined as: $\sigma_C(y) = (\sigma_j(y))_{j \in \mathbb{Z}_\nu^\times}$, where addition and multiplication in \mathbb{C}^n are component-wise. Indeed, for any $x, y \in K$, we have that $\sigma_C(x \cdot y)$ is the component wise product of $\sigma_C(x)$ and $\sigma_C(y)$. By elementary linear algebra, we observe that an element of K is fully specified by its canonical embedding vector.

The *trace* $\text{Tr} : K \rightarrow \mathbb{Q}$ and the (algebraic) *norm* $N : K \rightarrow \mathbb{Q}$ are defined as follows: $\text{Tr}(x) = \sum_{j \in \mathbb{Z}_\nu^\times} \sigma_j(x)$ and $N(x) = \prod_{j \in \mathbb{Z}_\nu^\times} \sigma_j(x)$. For any $x, y \in K$ we have $\text{Tr}(x \cdot y) = \sum_{j \in \mathbb{Z}_\nu^\times} \sigma_j(x) \cdot \sigma_j(y) = \langle \sigma_C(x), \overline{\sigma_C(y)} \rangle$ where $\langle \cdot, \cdot \rangle$ is the canonical Hermitian product on \mathbb{C}^n .

Space H . We use the following subspace of \mathbb{C}^n , as in [16]:

$$H = \{(x_j)_{j \in \mathbb{Z}_\nu^\times} \in \mathbb{C}^n : \forall j \in \mathbb{J}, x_{\nu-j} = \overline{x_j}\}.$$

Let $\mathbf{h}_j = \frac{1}{\sqrt{2}}(\mathbf{e}_j + \mathbf{e}_{\nu-j})$ and $\mathbf{h}_{\nu-j} = \frac{i}{\sqrt{2}}(\mathbf{e}_j - \mathbf{e}_{\nu-j})$ for $j \in \mathbb{J}$. The \mathbf{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 $(\mathbf{h}_j)_j$: For $x \in K$, we define $\sigma_H(x)$ by $\sigma_H(x) = (x_j)_j \in \mathbb{R}^n$ such that $\sigma_C(x) = \sum_j x_j \cdot \mathbf{h}_j$. As $\sigma_C(x) = (\sigma_j(x))_j$, we have, for $j \in \mathbb{J}$:

$$\begin{bmatrix} x_j \\ x_{\nu-j} \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ -i & i \end{bmatrix} \begin{bmatrix} \sigma_j(x) \\ \sigma_{\nu-j}(x) \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} \sigma_j(x) \\ \sigma_{\nu-j}(x) \end{bmatrix} = \begin{bmatrix} 1 & i \\ 1 & -i \end{bmatrix} \begin{bmatrix} x_j \\ x_{\nu-j} \end{bmatrix}.$$

The addition in H is component wise. Let $\sigma_H(x) = (x_j)_j$ and $\sigma_H(y) = (y_j)_j$, we define the multiplication by $\sigma_H(x \cdot y) = (z_j)_j$ where, for $j \in \mathbb{J}$:

$$\begin{bmatrix} z_j \\ z_{\nu-j} \end{bmatrix} = \begin{bmatrix} x_j & -x_{\nu-j} \\ x_{\nu-j} & x_j \end{bmatrix} \begin{bmatrix} y_j \\ y_{\nu-j} \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} z_j \\ z_{\nu-j} \end{bmatrix} = \begin{bmatrix} y_j & -y_{\nu-j} \\ y_{\nu-j} & y_j \end{bmatrix} \begin{bmatrix} x_j \\ x_{\nu-j} \end{bmatrix}.$$

To ease the presentation, we identify elements of K with their σ_H embeddings.

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 nonzero ideal, the *norm* $\mathcal{N}(I)$ of the ideal is the number of elements of the quotient ring R/I . We have $\mathcal{N}(x) = \mathcal{N}(I)$, 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 $I + J = \{\alpha + \beta : \alpha \in I, \beta \in J\}$. An ideal $I \subsetneq R$ is *prime* if for any $ab \in I$ then $a \in I$ or $b \in I$. Every ideal of R can be represented as a unique product of prime ideals, and for a prime ideal I , the quotient ring R/I is the finite field of order $\mathcal{N}(I)$. A *fractional ideal* $I \subseteq K$ is a set such that $dI \subseteq R$ is an (integral) ideal for a nonzero $d \in R$. 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 $I^\vee = \{x \in K : \text{Tr}(xI) \subseteq \mathbb{Z}\}$. We have $I^\vee = I^{-1} \cdot R^\vee$.

In our setup of cyclotomic fields, if q is a prime integer, the prime ideal factorization of $(q) \subseteq R$ can be computed efficiently. In particular, if $q = 1 \pmod{\nu}$, then $(q) = \prod_{j \in \mathbb{Z}_\nu^\times} \mathfrak{q}_j$ where each \mathfrak{q}_j is a prime ideal with norm $\mathcal{N}(\mathfrak{q}_j) = q$. The field K has n automorphisms $\tau_j : K \rightarrow K$ defined by $\tau_j(\xi) = \xi^j$ (for $j \in \mathbb{Z}_\nu^\times$). As noted in [16, Le. 2.16], the automorphism group of the $\{\tau_j\}$ acts transitively on the set $\{\mathfrak{q}_j\}_j$.

An isomorphism of quotient rings. Lyubashevsky et al. [16, Se. 2.3.9] used the Chinese Remainder Theorem to make 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^\vee respectively denote R/qR and R^\vee/qR^\vee .

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 [16, Le. 2.14]). The function $\theta_I : K \rightarrow K$ 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 $\theta_I^{-1} : I/qI \rightarrow R_q$ 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^\vee/qI^\vee to R_q^\vee that may be efficiently inverted using $\theta_I^{-1} : R_q^\vee \rightarrow I^\vee/qI^\vee$ with $\theta_I^{-1}(y) = t^{-1} \cdot y' \bmod qR$ where $y' = y \bmod qI^\vee$ and $y' \in (t)$.

Modules. A subset $M \subseteq K^d$ is an R -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 $(\mathbf{b}_k)_k$ of vectors in K^d such that $M = \sum_k R \cdot \mathbf{b}_k$. 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., [9, Ch. 1]): For every module M , there exist I_k nonzero ideals of R and $(\mathbf{b}_k)_k$ linearly independent vectors of K^d such that $M = \sum_{k=1}^d I_k \cdot \mathbf{b}_k$. We say that $[(I_k)_k, (\mathbf{b}_k)_k]$ is a *pseudo-basis* of M . The terminology pseudo-basis is used as the coefficient ideals I_k can be non-principal. The representation of the elements of M with respect to a pseudo-basis is unique. Two pseudo-bases can generate the same module and then, they have the same cardinality. The latter is called *rank* of the module.

We define the dual of a module by $M^\vee = \{\mathbf{x} \in K^d, \forall \mathbf{y} \in M : \text{Tr}(\langle \mathbf{x}, \overline{\mathbf{y}} \rangle) \in \mathbb{Z}\}$, where $\langle \cdot, \cdot \rangle$ is the Hermitian product on K^d . We have the following property:

Lemma 2.1. *If $M = \sum_{k=1}^d I_k \cdot \mathbf{b}_k$, then $M^\vee = \sum_{k=1}^d I_k^\vee \cdot \mathbf{b}_k^\vee$, where the \mathbf{b}_ℓ^\vee 's are defined by $\langle \mathbf{b}_k, \overline{\mathbf{b}_\ell^\vee} \rangle = 1$ if $k = \ell$ and $\langle \mathbf{b}_k, \overline{\mathbf{b}_\ell^\vee} \rangle = 0$ otherwise.*

Proof. We first show that $\sum_{k=1}^d I_k^\vee \cdot \mathbf{b}_k^\vee \subseteq M^\vee$. Let $\mathbf{x} \in \sum_{k=1}^d I_k^\vee \cdot \mathbf{b}_k^\vee$. Then for each i there exists $x_k \in I_k^\vee$ such that $\mathbf{x} = \sum_{k=1}^d x_k \cdot \mathbf{b}_k^\vee$. Let $\mathbf{y} = \sum_{k=1}^d y_k \cdot \mathbf{b}_k \in M$. Then by linearity, we have $\text{Tr}(\langle \mathbf{x}, \overline{\mathbf{y}} \rangle) = \sum_{k=1}^d \text{Tr}(x_k y_k)$. For all i , we have $x_k \in I_k^\vee$ and $y_k \in I_k$, and thus $\text{Tr}(x_k y_k) \in \mathbb{Z}$. Therefore, we have $\text{Tr}(\langle \mathbf{x}, \overline{\mathbf{y}} \rangle) \in \mathbb{Z}$ and $\mathbf{x} \in M^\vee$.

We now show that $M^\vee \subseteq \sum_{k=1}^d I_k^\vee \cdot \mathbf{b}_k^\vee$. Let $\mathbf{x} \in M^\vee \subseteq K^d$. We can write $\mathbf{x} = \sum_{k=1}^d x_k \cdot \mathbf{b}_k^\vee$, for some x_k 's in K . It suffices to show that $x_k \in I_k^\vee$. Let $y_k \in I_k$ be arbitrary. By linearity, we have $\text{Tr}(\langle \mathbf{x}, \overline{y_k \mathbf{b}_k} \rangle) = \text{Tr}(x_k y_k) \in \mathbb{Z}$. This implies that $x_k \in I_k^\vee$. \square

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

2.2 Lattices

We refer to [21,34] for introductions to lattices and their computational aspects. A *euclidean lattice* $\Lambda \subseteq \mathbb{R}^n$ is the set of all integer linear combinations $\sum_{i=1}^p x_i \mathbf{b}_i$ of some linearly independent vectors $(\mathbf{b}_i)_{1 \leq i \leq p} \in \mathbb{R}^n$. We write $\mathcal{L}(\mathbf{B})$ for the lattice spanned by the basis $\mathbf{B} = (\mathbf{b}_i)_{i \leq p}$. 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 nonzero vectors. More generally, the *i th 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 $\Lambda^* = \{\mathbf{x} \in \mathbb{R}^n : \forall \mathbf{y} \in \Lambda, \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$. If $\Lambda = \mathcal{L}(\mathbf{B})$ then $\Lambda^* = \mathcal{L}(\mathbf{B}^*)$ with $\mathbf{B}^* = \mathbf{B}^{-T}$.

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 $_\gamma$). Let $\gamma \geq 1$ be a function of the dimension n . The *Generalized Independent Vectors Problem* GIVP $_\gamma^\phi$ is as follows: Given a lattice basis \mathbf{B} , find $n = \dim(\mathcal{L}(\mathbf{B}))$ linearly independent vectors $\mathbf{s}_1, \dots, \mathbf{s}_n \in \mathcal{L}(\mathbf{B})$ such that $\max_i \|\mathbf{s}_i\| \leq \gamma \cdot \phi(\mathcal{L}(\mathbf{B}))$.

For $\phi = \lambda_n$, this problem is NP-hard for any approximation factor $\gamma \leq O(1)$ (see [4]). The best known algorithms (even quantum) for an exact solution and an approximation to within any polynomial factor γ all have exponential complexities [25]. 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.

Ideal and module lattices. As σ_H is an embedding from K to \mathbb{R}^n and I an ideal of R , the set $\sigma_H(I)$ is a lattice. We call it ideal lattice with respect to K . To ease the presentation, we often identify I and $\sigma_H(I)$. We let Id-GIVP denote the restriction of GIVP to ideal lattices.

We define module lattices similarly. The map $(\sigma_H, \dots, \sigma_H)$ is an embedding from K^d to \mathbb{R}^N , with $N = nd$, and $M \subseteq K^d$ a module of R . By abuse of notation, we also call it σ_H . The set $\sigma_H(M)$ is a module lattice. Similarly to ideal lattices, we let Mod-GIVP denote the restriction of GIVP to module lattices. Note that if M is a rank d module and if K has degree n , then the corresponding module lattice has dimension $N = nd$. For any $\mathbf{x} \in K^d$, we define $\|\mathbf{x}\| = (\sum_{k \in [d]} \sum_{j \in \mathbb{Z}_\nu^\times} |\sigma_j(x_k)|^2)^{1/2}$. We also define $\|\mathbf{x}\|_\infty = \max_{j,k} |\sigma_j(x_k)|$, $\|\mathbf{x}\|_{2,\infty} = \max_j (\sum_k |\sigma_j(x_k)|^2)^{1/2}$ and $\|\mathbf{x}\|_{\infty,2} = \max_k (\sum_j |\sigma_j(x_k)|^2)^{1/2}$.

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 [6,9]. All asymptotic statements involving modules (including hardness results) will be given for N growing to infinity.

2.3 Gaussian measures

For a vector $\mathbf{c} \in \mathbb{R}^n$ and a real $s > 0$, the Gaussian function is defined by $\rho_{s,\mathbf{c}}(\mathbf{x}) = e^{-\pi \|\frac{\mathbf{x}-\mathbf{c}}{s}\|^2}$, for all $\mathbf{x} \in \mathbb{R}^n$. This function is extended to any countable set $A \subseteq \mathbb{R}^n$ in the usual way: $\rho_{s,\mathbf{c}}(A) = \sum_{\mathbf{x} \in A} \rho_{s,\mathbf{c}}(\mathbf{x})$. By normalizing the Gaussian function, we obtain the continuous Gaussian probability distribution: $D_s(\mathbf{x}) = \rho_s(\mathbf{x})/s^n$. For $\mathbf{r} = (r_1, \dots, r_n)^T \in (\mathbb{R}^+)^n$, a sample from $D_{\mathbf{r}}$ over \mathbb{R}^n is given by $(D_{r_i})_i$.

In the following of the paper, we use the elliptical Gaussian distributions in the basis $\{\mathbf{h}_j\}_{j \in \mathbb{Z}_\nu^\times}$, as in [16]. For $(r_j)_{j \in \mathbb{Z}_\nu^\times} \in \mathbb{R}^n$ such that $r_j = r_{\nu-j}$ for all $j \in \mathbb{J}$, a sample x from $D_{\mathbf{r}}$ is given by $\sigma_C(x) = \sum_j x_j \cdot \mathbf{h}_j$, where each x_j is independently chosen from the Gaussian distribution D_{r_j} over \mathbb{R} .

We define $\Psi_{[\alpha,\alpha']}$ for $0 \leq \alpha < \alpha'$, as the set of Gaussian distributions $D_{\mathbf{r}}$ with $\alpha < r_i \leq \alpha'$, for all i . We write $\Psi_{\leq \alpha'}$ when $\alpha = 0$. We also recall the distribution Υ_α used in [16]. The gamma distribution $\Gamma(2, 1)$ with shape parameter 2 and scale parameter 1 has density $t \exp(-t)$ for $t \geq 0$ and zero for $t < 0$. For $\alpha > 0$, a distribution sampled from Υ_α is an elliptical Gaussian distribution $D_{\mathbf{r}}$ whose parameters are $r_j = r_{\nu-j} = \alpha \sqrt{1 + \sqrt{n} x_j}$, where the x_j 's for $j \in \mathbb{J}$ are chosen independently from $\Gamma(2, 1)$. We will use the following result on $\Gamma(2, 1)$.

Lemma 2.2 ([16, Claim 5.10]). *Let P be the distribution $\Gamma(2, 1)^n$ and Q be the distribution $(\Gamma(2, 1) - z_1) \times \dots \times (\Gamma(2, 1) - z_n)$ for some $0 \leq z_1, \dots, z_n \leq 1/\sqrt{n}$. Then for any measurable set $A \subseteq \mathbb{R}^n$, we have $\int_A Q \geq \frac{1}{\text{poly}(n)} \cdot (\int_A P)^2$.*

For all $\mathbf{c} \in \mathbb{R}^n$, $s > 0$ and lattice Λ , the discrete Gaussian probability distribution with support Λ , center \mathbf{c} and standard deviation s is defined by:

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

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

Theorem 2.3 ([11, Th. 4.1] and [8, Se. 5]). *There is a probabilistic polynomial time algorithm that, given a basis \mathbf{B} of an n -dimensional lattice $\Lambda = \mathcal{L}(\mathbf{B})$, a standard deviation $s \geq \|\tilde{\mathbf{B}}\| \cdot \sqrt{\log n}$, and a center $\mathbf{c} \in \mathbb{R}^n$, outputs a sample whose distribution is $D_{\Lambda,s,\mathbf{c}}$.*

The *smoothing parameter* of a lattice was introduced by [24]. 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 2.4 ([24, Le. 3.3]). *Let Λ be an n -dimensional lattice and $\varepsilon > 0$. Then $\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}} \cdot \lambda_n(\Lambda)$.*

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

Lemma 2.5 ([28, 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 lattice minimum with respect to the infinity norm.*

Lemma 2.6 ([11, 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 $\mathbf{c} \in \mathbb{R}^n$, the distribution $(D_{\Lambda, s, \mathbf{c}} \bmod \Lambda')$ is within statistical distance at most 2ε of the uniform distribution over Λ/Λ' .*

Lemma 2.7 ([23, Le. 4.4]). *Let Λ be an n -dimensional lattice, $s > 2\eta_\varepsilon(\Lambda)$ for $\varepsilon \leq 1/100$, and $\mathbf{c} \in \mathbb{R}^n$. Then for any $(n-1)$ -dimensional hyperplane \mathcal{H} , the probability that $x \notin \mathcal{H}$ where x is chosen from $D_{\Lambda, s, \mathbf{c}}$ 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 similar result for the sum of a continuous Gaussian and a discrete one.

Lemma 2.8 (Adapted from [36, Claim 3.9]). *Let Λ be an n -dimensional lattice, $\mathbf{u} \in \mathbb{R}^n$, $\mathbf{r} \in (\mathbb{R}^+)^n$, $\sigma > 0$ and $t_i = \sqrt{r_i^2 + \sigma^2}$ for all i . Assume that $\min_i r_i \sigma / t_i \geq \eta_\varepsilon(\Lambda)$ for some $\varepsilon \in (0, 1/2)$. Consider the discrete distribution Y on \mathbb{R}^n obtained by sampling from $D_{\Lambda + \mathbf{u}, \mathbf{r}}$ and then adding a vector taken from D_σ . Then we have $\Delta(Y, D_t) \leq 4\varepsilon$.*

Proof. This proof follows the same principle as the one of [36, Claim 3.9], the only difference being that [36, Claim 3.9] considers the case where all r_i 's are equal. Using the Poisson summation formula, one obtains that the probability density function Y can be written as:

$$\forall \mathbf{x} \in \mathbb{R}^n : Y(\mathbf{x}) = \frac{\rho_{\mathbf{t}}(\mathbf{x})}{\prod_i t_i} \cdot \frac{\left(\prod_i \frac{t_i}{\sigma r_i}\right) \cdot \widehat{\rho_{\mathbf{t}', \mathbf{x}' - \mathbf{u}}}(\Lambda^*)}{\left(\prod_i \frac{1}{r_i}\right) \cdot \widehat{\rho_{\mathbf{r}, -\mathbf{u}}}(\Lambda^*)},$$

where $t'_i = r_i \sigma / t_i$ and $x'_i = r_i^2 x_i / t_i^2$ for all i , and where \widehat{f} denotes the Fourier transform of f . Then, we have:

$$\left| 1 - \left(\prod_i \frac{t_i}{\sigma r_i}\right) \widehat{\rho_{\mathbf{t}', \mathbf{x}' - \mathbf{u}}}(\Lambda^*) \right| \leq \rho_{\mathbf{t}''}(\Lambda^* \setminus \{0\}), \quad \text{with } t''_i = 1/t'_i \text{ for all } i,$$

$$\left| 1 - \left(\prod_i \frac{1}{r_i}\right) \widehat{\rho_{\mathbf{r}, -\mathbf{u}}}(\Lambda^*) \right| \leq \rho_{\mathbf{r}''}(\Lambda^* \setminus \{0\}), \quad \text{with } r''_i = 1/r_i \text{ for all } i.$$

Let \mathbf{s}' and $\sigma' > 0$ be such that $s'_i \geq \sigma'$ for all i . We have that for any vector \mathbf{x} :

$$\frac{\rho_{1/\sigma'}(\mathbf{x})}{\rho_{(1/s'_i)_i}(\mathbf{x})} = \exp\left(-\pi \sum_i ((\sigma')^2 x_i^2 - (s'_i)^2 x_i^2)\right) \geq 1.$$

This implies that $\rho_{\mathbf{t}''}(\Lambda^* \setminus \{0\}) \leq \varepsilon$ and $\rho_{\mathbf{r}''}(\Lambda^* \setminus \{0\}) \leq \varepsilon$, which completes the proof. \square

We will also use the following lemmata.

Lemma 2.9 (Adapted from [28, Cor. 5.3]). For any n -dimensional lattice $\Lambda \subseteq \mathbb{R}^n$, $\mathbf{c} \in \mathbb{R}^n$, $\varepsilon \in (0, 1)$, $t \geq \sqrt{2\pi}$, unit vector $\mathbf{u} \in \mathbb{R}^n$ and $s \geq \eta_\varepsilon(\Lambda)$, we have:

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

Lemma 2.10 (Adapted from [28, Cor. 5.3]). Let Λ be a n -dimensional lattice, $\varepsilon \in (0, 1)$ and $\mathbf{r} \in \mathbb{R}^n$ with $r_i \geq \eta_\varepsilon(\Lambda)$ for all $i \leq n$. Then we have $\Pr_{\mathbf{x} \leftarrow D_{\Lambda, \mathbf{r}}} [\|\mathbf{x}\|_\infty \geq (\max_i r_i) \cdot t] \leq 2en \cdot \exp(-\pi t^2)$ for all $t > 0$. In particular, for $t = \omega(\sqrt{\log n})$ (resp. $t = \Omega(\sqrt{n})$) the above probability is at most $n^{-\omega(1)}$ (resp. $2^{-\Omega(n)}$).

We now generalize [28, Cor. 5.3] and [37, Le. 2.9] to the case of module lattices (over the ring of integers of a cyclotomic number field).

Lemma 2.11. Let $\varepsilon \in (0, \frac{1}{2m+1})$ and $z_1, \dots, z_m \in R$. Let $M \subseteq K^d$ be a rank d module on R , $s \geq \eta_\varepsilon(M)$ and $\mathbf{c}_1, \dots, \mathbf{c}_m \in R^d$. If the \mathbf{y}_ℓ 's are independently sampled from the $D_{M, s, \mathbf{c}_\ell}$'s, then, for all $t \geq 0$:

$$\Pr \left[\left\| \sum_{\ell \in [m]} z_\ell \cdot (\mathbf{y}_\ell - \mathbf{c}_\ell) \right\|_\infty \geq st \|\mathbf{z}\| \right] \leq 2 \frac{1 + \varepsilon}{1 - \varepsilon} t N \sqrt{2\pi e} \cdot e^{-\pi t^2}.$$

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

Proof. The proof builds upon that of [28, Cor. 5.3]. The principle is to interpret the m Gaussian samples from the N -dimensional lattice M as one Gaussian sample from the (Nm) -dimensional lattice L and then apply Lemma 2.9, where $L = M \times \dots \times M$ (i.e., the Cartesian product of m copies of M). We also define $\vec{\mathbf{c}} = (\mathbf{c}_1, \dots, \mathbf{c}_m)^T \in (R^d)^m$ and $\vec{\mathbf{y}} = (\mathbf{y}_1, \dots, \mathbf{y}_m)^T \in (R^d)^m$. We have $\rho_{s, \vec{\mathbf{c}}}(L) = \prod_{\ell \in [m]} \rho_{s, \mathbf{c}_\ell}(M)$. The vector $\vec{\mathbf{y}}$ has distribution $D_{L, s, \vec{\mathbf{c}}}$. We have:

$$\sigma_C \left(\sum_{\ell \in [m]} z_\ell \cdot (\mathbf{y}_\ell - \mathbf{c}_\ell) \right) = \begin{bmatrix} \sum_{\ell=1}^m \sigma_C(z_\ell \cdot (y_\ell^{(1)} - c_\ell^{(1)})) \\ \vdots \\ \sum_{\ell=1}^m \sigma_C(z_\ell \cdot (y_\ell^{(d)} - c_\ell^{(d)})) \end{bmatrix} = \begin{bmatrix} \left(\langle \sigma_j(\vec{\mathbf{z}}), \overline{\sigma_j(\vec{\mathbf{y}}^{(1)} - \vec{\mathbf{c}}^{(1)})} \rangle \right)_{j \in \mathbb{Z}_\nu^\times} \\ \vdots \\ \left(\langle \sigma_j(\vec{\mathbf{z}}), \overline{\sigma_j(\vec{\mathbf{y}}^{(d)} - \vec{\mathbf{c}}^{(d)})} \rangle \right)_{j \in \mathbb{Z}_\nu^\times} \end{bmatrix}$$

with $\vec{\mathbf{z}} = (z_1, \dots, z_m)^T \in R^m$, $\vec{\mathbf{y}}^{(k)} - \vec{\mathbf{c}}^{(k)} = (y_1^{(k)} - c_1^{(k)}, \dots, y_m^{(k)} - c_m^{(k)})^T \in R^m$ for $k \in [d]$, and, for any $j \in \mathbb{Z}_\nu^\times$ and $\vec{\mathbf{x}} \in R^m$, $\sigma_j(\vec{\mathbf{x}}) = (\sigma_j(x_\ell))_{\ell \in [m]}$.

By applying the union bound over all $j \in \mathbb{Z}_\nu^\times$ and all $k \in [d]$, it suffices to obtain a probabilistic upper bound on the Hermitian product between $\sigma_j(\vec{\mathbf{z}})$ and $\overline{\sigma_j(\vec{\mathbf{y}}^{(k)} - \vec{\mathbf{c}}^{(k)})}$ for any fixed j and k . For the rest of the proof, we fix $j \in \mathbb{Z}_\nu^\times$ and $k \in [d]$. Wlog (by complex conjugation), we take $j \in \mathbb{J}$.

For $\ell \in [m]$, let $\mathbf{u}_\ell = (u_\ell^{(1)}, \dots, u_\ell^{(d)})^T \in \mathbb{C}^{nd}$ with $u_\ell^{(k')} = (0, \dots, 0)^T$ for $k' \neq k$, and:

$$u_\ell^{(k)} = (0, \dots, 0, \sigma_j(z_\ell), 0, \dots, 0, -i \cdot \sigma_j(z_\ell), 0, \dots, 0)^T,$$

i.e., the coordinate of index j is equal to $\sigma_j(z_\ell)$, the coordinate of index $\nu - j$ is equal to $-i \cdot \sigma_j(z_\ell)$, and all the others are 0. We now define $\vec{\mathbf{u}} \in \mathbb{C}^{ndm}$ as the concatenation of the \mathbf{u}_ℓ 's (for $\ell \in [m]$), and $\sigma_H(\vec{\mathbf{y}} - \vec{\mathbf{c}}) \in \mathbb{R}^{ndm}$ as the concatenation of the $\sigma_H(\mathbf{y}_\ell - \mathbf{c}_\ell)$'s. We have:

$$\langle \sigma_j(\vec{\mathbf{z}}), \overline{\sigma_j(\vec{\mathbf{y}}^{(k)} - \vec{\mathbf{c}}^{(k)})} \rangle = \sum_{\ell} \sigma_j(z_\ell) \overline{\sigma_j(y_\ell^{(k)} - c_\ell^{(k)})} = \langle \vec{\mathbf{u}}, \sigma_H(\vec{\mathbf{y}} - \vec{\mathbf{c}}) \rangle.$$

Now, we define $\vec{\mathbf{v}} = \vec{\mathbf{u}} / \|\vec{\mathbf{u}}\| \in \mathbb{C}^{ndm}$. By Lemma 2.9, we have:

$$\begin{aligned} \Pr_{\vec{\mathbf{y}} \leftarrow D_{L, s, \vec{\mathbf{c}}}} [|\langle \sigma_H(\vec{\mathbf{y}} - \vec{\mathbf{c}}), \Re(\vec{\mathbf{v}}) \rangle| \geq st] &\leq \frac{1 + \varepsilon}{1 - \varepsilon} t \sqrt{2\pi e} \cdot e^{-\pi t^2}, \\ \Pr_{\vec{\mathbf{y}} \leftarrow D_{L, s, \vec{\mathbf{c}}}} [|\langle \sigma_H(\vec{\mathbf{y}} - \vec{\mathbf{c}}), \Im(\vec{\mathbf{v}}) \rangle| \geq st] &\leq \frac{1 + \varepsilon}{1 - \varepsilon} t \sqrt{2\pi e} \cdot e^{-\pi t^2}, \end{aligned}$$

where \Re and \Im respectively denote the real and imaginary parts of a complex number. By using the union bound and scaling by $\|\vec{\mathbf{u}}\| \leq \|\mathbf{z}\|$, we obtain that:

$$\Pr [|\langle \vec{\mathbf{u}}, \sigma_H(\vec{\mathbf{y}}' - \vec{\mathbf{c}}') \rangle| \geq st\|\mathbf{z}\|] \leq 2 \frac{1+\varepsilon}{1-\varepsilon} t \sqrt{2\pi e} \cdot e^{-\pi t^2}.$$

This leads to the claimed result. \square

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 [16], but without proof.

Lemma 2.12. *Let $\mathbf{r} \in (\mathbb{R}^+)^n$ with $r_j = r_{\nu-j}$ for all $j \in \mathbb{Z}_\nu^\times$, $x \in K$ sampled from $D_{\mathbf{r}}$ and $e \in K$ fixed. Then $x \cdot e$ is distributed from $D_{\mathbf{r}'}$ with $r'_j = r_j |\sigma_j(e)|$ for all j .*

Proof. Let us write $\sigma_C(x) = \sum_j x_j \cdot \mathbf{h}_j$ where each x_j is sampled from D_{r_j} . By definition of the \mathbf{h}_j 's, we have $\sigma_j(x) = (x_j + ix_{\nu-j})$ and $\sigma_{\nu-j}(x) = (x_j - ix_{\nu-j})$, for $j \in \mathbb{J}$. Let $\sigma_C(e) = \sum_j e_j \cdot \mathbf{h}_j$ and $\sigma_C(e \cdot x) = \sum_j y_j \cdot \mathbf{h}_j$. We have, for $j \in \mathbb{J}$

$$\begin{bmatrix} y_j \\ y_{\nu-j} \end{bmatrix} = \begin{bmatrix} e_j & -e_{\nu-j} \\ e_{\nu-j} & e_j \end{bmatrix} \begin{bmatrix} x_j \\ x_{\nu-j} \end{bmatrix}$$

The vector $(y_j, y_{\nu-j})^T$ is an orthogonal transformation of the vector $(x_j, x_{\nu-j})$, and thus y_j and $y_{\nu-j}$ are statistically independent. Further, the reals y_j and $y_{\nu-j}$ are samples of $D_{r'_j}$ and $D_{r'_{\nu-j}}$ respectively, with $r'_j = r'_{\nu-j} = (e_j^2 r_j^2 + e_{\nu-j}^2 r_{\nu-j}^2)^{1/2} = r_j |\sigma_j(e)|$. \square

The following lemma generalizes the previous result to the module setting.

Lemma 2.13. *Let $\mathbf{r} \in (\mathbb{R}^+)^n$ with $r_j = r_{\nu-j}$ for all $j \in \mathbb{Z}_\nu^\times$, $\mathbf{x} \in K^d$ sampled from $D_{\mathbf{s}, \dots, \mathbf{s}}$ and $\mathbf{e} \in K^d$ fixed. Then $\sum_k x_k e_k$ is distributed from $D_{\mathbf{r}'}$ with $r'_j = r_j \cdot (\sum_{k \in [d]} |\sigma_j(e_k)|^2)^{1/2}$ for all j .*

Proof. By Lemma 2.12, we have that $x_k \cdot e_k$ has distribution $D_{\mathbf{r}'_k}$ with $r'_{k,j} = r'_{k,\nu-j} = r_j |\sigma_j(e_k)|$ for all j . The quantity under scope is the sum of independent Gaussians. \square

3 The Short Integer Solution Problem

In this section, we describe a reduction from Mod-GIVP to M-SIS.

3.1 Variants of SIS

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

The Short Integer Solution Problem. The SIS problem was first introduced by Ajtai [1].

Definition 3.1. *The Small Integer Solution problem $\text{SIS}_{q,m,\beta}$ is as follows: Given $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ chosen from the uniform distribution, find $\mathbf{z} \in \mathbb{Z}^m$ such that $\mathbf{Az} = 0 \pmod q$ and $0 < \|\mathbf{z}\| \leq \beta$.*

As observed in [24, Le. 5.2], for any q , $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\beta \geq \sqrt{mq}^{n/m}$, the SIS instance (q, \mathbf{A}, β) admits a solution. There are several reductions from GIVP to SIS (see, e.g., [1,23,11]). The strongest known result is the following.

Theorem 3.2 (Adapted from [11, Th. 9.2]). *For $\varepsilon(n) = n^{-\omega(1)}$, there is a probabilistic polynomial time reduction from solving $\text{GIVP}_\gamma^{\eta_\varepsilon}$ in polynomial time (in the worst case, with high probability) to solving $\text{SIS}_{q,m,\beta}$ in polynomial time with non-negligible probability, for any $m(n), q(n), \beta(n)$ and $\gamma(n)$ such that $\gamma \geq \beta \sqrt{n} \cdot \omega(\sqrt{\log n})$, $q \geq \beta \sqrt{n} \cdot \omega(\log n)$ and $m, \log q \leq \text{poly}(n)$.*

SIS over rings. The R-SIS problem was concurrently introduced in [31] and [15].

Definition 3.3. *The problem R-SIS $_{q,m,\beta}$ is as follows: Given $a_1, \dots, a_m \in R_q$ chosen independently from the uniform distribution, find $z_1, \dots, z_m \in R$ such that $\sum_{i=1}^m a_i \cdot z_i = 0 \pmod q$ and $0 < \|\mathbf{z}\| \leq \beta$, where $\mathbf{z} = (z_1, \dots, z_m)^T \in R^m$.*

This problem over rings can be interpreted in terms of structured integer matrices. For example, when n is a power of 2, then R and R_q are isomorphic to $\mathbb{Z}[x]/(x^n + 1)$ and $\mathbb{Z}_q[x]/(x^n + 1)$ respectively, and the ring multiplication $a_i \cdot z_i$ can be written as the multiplication of the vector of \mathbb{Z}^n whose entries are the coefficients of z_i and, with a nega-circulant matrix whose entries are derived from the coefficients of a_i . In this setup, R-SIS is a variant of SIS where \mathbf{A} is restricted to being block negacirculant: $\mathbf{A} = [\text{Rot}(a_1) | \dots | \text{Rot}(a_m)]$, with:

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

By using a technique from [16] (namely, the isomorphism between I/qI and R_q described in Subsection 2.1) into the proof of [15], one obtains the following result.

Theorem 3.4. *For $\varepsilon(n) = n^{-\omega(1)}$, there is a probabilistic polynomial time reduction from solving Id-GIVP $_{\gamma}^{\eta\varepsilon}$ in polynomial time (in the worst case, with high probability) to solving R-SIS $_{q,m,\beta}$ in polynomial time with non-negligible probability, for any $m(n), q(n), \beta(n)$ and $\gamma(n)$ such that $\gamma \geq \beta\sqrt{n} \cdot \omega(\sqrt{\log n})$, $q \geq \beta\sqrt{n} \cdot \omega(\log n)$ and $m, \log q \leq \text{poly}(n)$.*

SIS over modules. The problem M-SIS generalizes both SIS and R-SIS. We use the following notations: the variable n denotes the dimension of the ring R and the variable d corresponds to the rank of the module $M \subseteq R^d$; we let $N = nd$ denote the dimension of the corresponding module lattice, and give the complexity statements for N growing to infinity.

Definition 3.5. *The problem M-SIS $_{q,m,\beta}$ is as follows: Given $\mathbf{a}_1, \dots, \mathbf{a}_m \in R_q^d$ chosen independently from the uniform distribution, find $z_1, \dots, z_m \in R$ such that $\sum_{i=1}^m \mathbf{a}_i \cdot z_i = 0 \pmod q$ and $0 < \|\mathbf{z}\| \leq \beta$, where $\mathbf{z} = (z_1, \dots, z_m)^T \in R^m$.*

Like R-SIS, M-SIS can be interpreted in terms of matrices. In the same setting as above for R-SIS, it consists in taking a SIS matrix \mathbf{A} of the form:

$$\begin{array}{c} \begin{array}{c} \xrightarrow{n = N/d} \\ \left[\begin{array}{cc|ccc|cc} \text{Rot}(\mathbf{a}_{1,1}) & \text{Rot}(\mathbf{a}_{1,2}) & \cdots & \text{Rot}(\mathbf{a}_{1,m-1}) & \text{Rot}(\mathbf{a}_{1,m}) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \text{Rot}(\mathbf{a}_{d,1}) & \text{Rot}(\mathbf{a}_{d,2}) & \cdots & \text{Rot}(\mathbf{a}_{d,m-1}) & \text{Rot}(\mathbf{a}_{d,m}) \end{array} \right] \\ \xrightarrow{m \text{ blocks}} \end{array} \\ \begin{array}{c} \left. \begin{array}{c} \vdots \\ \vdots \\ \vdots \end{array} \right\} d \text{ blocks} \end{array} \end{array}$$

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

Theorem 3.6. *For any $d \geq 1$ and $\varepsilon(N) = N^{-\omega(1)}$, there is a probabilistic polynomial time reduction from solving Mod-GIVP $_{\gamma}^{\eta\varepsilon}$ in polynomial time (in the worst case, with high probability) to solving M-SIS $_{q,m,\beta}$ in polynomial time with non-negligible probability, for any $m(N), q(N), \beta(N)$ and $\gamma(N)$ such that $\gamma \geq \beta\sqrt{N} \cdot \omega(\sqrt{\log N})$, $q \geq \beta\sqrt{N} \cdot \omega(\log N)$ and $m, \log q \leq \text{poly}(N)$.*

In the case of a sub-exponential oracle (and with $\varepsilon(N) = 2^{-\Omega(N)}$), the result still holds and the conditions on the parameters become $\gamma \geq \beta\Omega(N)$ and $q \geq \beta\Omega(N^{3/2})$.

Taking $n = N$ and $d = 1$ in Theorem 3.6 allows us to recover Theorem 3.4. Also, by taking $n = 1$ and $d = N$ in Theorem 3.6, we obtain a hardness result for SIS that is as good as that of Theorem 3.2.

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 [23].

Definition 3.7 ([23, Def. 5.3]). *The Incremental Independent Vectors Problem $\text{IncGIVP}_\gamma^{\eta_\varepsilon}$, is as follows: Given a tuple $(\mathbf{B}, \mathbf{S}, \mathcal{H})$ where \mathbf{B} is a basis of an n -dimensional lattice, $\mathbf{S} \subseteq \mathcal{L}(\mathbf{B})$ is a full-rank set of vectors such that $\|\mathbf{S}\| \geq \gamma \cdot \eta_\varepsilon(\mathcal{L}(\mathbf{B}))$ and \mathcal{H} is a hyperplane, find $\mathbf{h} \in \mathcal{L}(\mathbf{B}) \setminus \mathcal{H}$ such that $\|\mathbf{h}\| \leq \|\mathbf{S}\|/2$.*

Theorem 3.8 ([19, Th. 6.3]). *For any function ε and γ , there is a probabilistic polynomial time reduction from solving $\text{GIVP}_\gamma^{\eta_\varepsilon}$ (in the worst case, with high probability) to solving $\text{IncGIVP}_\gamma^{\eta_\varepsilon}$ (in the worst case, with high probability).*

As the latter reduction preserves the lattice, it induces a reduction from $\text{Mod-GIVP}_\gamma^{\eta_\varepsilon}$ to $\text{Mod-IncGIVP}_\gamma^{\eta_\varepsilon}$, i.e., $\text{IncGIVP}_\gamma^{\eta_\varepsilon}$ restricted to module lattices. To prove Theorem 3.6, we provide a reduction from $\text{Mod-IncGIVP}_\gamma^{\eta_\varepsilon}$ to $\text{M-SIS}_{q,m,\beta}$.

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

$$\max\left(\frac{2q}{\gamma}, \sqrt{\log N}\right) \|\mathbf{S}\| \leq s \leq \frac{q\|\mathbf{S}\|}{2\beta\sqrt{N} \cdot \omega(\sqrt{\log N})}$$

- For all $\ell \leq m$,
 - Get a fresh \mathbf{y}_ℓ distributed as $D_{\mathcal{L}(\mathbf{B}),s,\mathbf{0}}$ (using Theorem 2.3),
 - Let $\mathbf{a}_\ell = \Theta^{-1}(\mathbf{y}_\ell \bmod qM)$ (see the definition of Θ in Section 2.1).
- Invoke the oracle \mathcal{O} on input $(\mathbf{a}_1, \dots, \mathbf{a}_m)$. If \mathcal{O} succeeds, it returns $\mathbf{z} = (z_1, \dots, z_m)^T \in R^m$ such that $\sum_{\ell \in [m]} \mathbf{a}_\ell \cdot z_\ell = \mathbf{0} \bmod q$ and $0 < \|\mathbf{z}\| \leq \beta$.
- Output $\mathbf{h} = \frac{1}{q} \sum_{\ell \in [m]} z_\ell \cdot \mathbf{y}_\ell$.

This algorithm runs in polynomial time. Also, thanks to the parameter constraints, the interval to which the standard deviation s must belong is nonempty. Moreover, the standard deviation s is sufficiently large for the assumptions of Theorem 2.3 to hold. Indeed, by [21, 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 $\|\mathbf{T}\| \leq \|\mathbf{S}\| \leq \|\mathbf{S}\|$. We use this basis and we have that $s \geq \|\mathbf{T}\| \cdot \sqrt{\log N}$.

Lemma 3.9. *The statistical distance between the distribution of $(\mathbf{a}_1, \dots, \mathbf{a}_m)$ and the uniform distribution over R_q^d is at most $2m\varepsilon$.*

Proof. We have $s \geq \frac{2q}{\gamma} \cdot \|\mathbf{S}\|$ and $\|\mathbf{S}\| \geq \gamma \cdot \eta_\varepsilon(M)$. This implies that $s \geq q \cdot \eta_\varepsilon(M) = \eta_\varepsilon(qM)$. By Lemma 2.6 applied to the lattices M and qM , the statistical distance between the distribution of $(\mathbf{y}_\ell \bmod qM)$ 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 $\mathbf{a}_\ell = \Theta^{-1}(\mathbf{y}_\ell)$ and the uniform distribution on $(R/qR)^d$ is also at most 2ε . The result follows. \square

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

Lemma 3.10. *For any hyperplane \mathcal{H} , the probability that the output vector \mathbf{h} does not belong to \mathcal{H} is $\geq 1/100$.*

Proof. As \mathcal{O} succeeded, the vector \mathbf{z} is nonzero. By definition of \mathbf{h} , for every \mathbf{y}'_1 we have:

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

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

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

As $s \geq 2q \cdot \eta_\varepsilon(M) = 2\eta_\varepsilon(qM)$, Lemma 2.7 gives that this probability is $\geq 1/100$. \square

The following completes the proof of Theorem 3.6.

Lemma 3.11. *We have $\mathbf{h} \in M$ and, with probability close to 1, we have that $\|\mathbf{h}\| \leq \|\mathbf{S}\|/2$.*

Proof. Let us first show that $\mathbf{h} \in M$. We have, modulo qM :

$$\sum_{\ell=1}^m z_\ell \cdot \mathbf{y}_\ell = \sum_{\ell=1}^m z_\ell \cdot \Theta(\mathbf{a}_\ell) = \Theta\left(\sum_{\ell=1}^m z_\ell \mathbf{a}_\ell\right) = \mathbf{0}.$$

This implies that $\mathbf{h} = (\sum_{\ell=1}^m z_\ell \cdot \mathbf{y}_\ell)/q$ belongs to M .

We now show that $\|\mathbf{h}\| \leq \|\mathbf{S}\|/2$. We have $\|\mathbf{h}\| = \|\sum_{\ell=1}^m z_\ell \cdot \mathbf{y}_\ell\|/q$. As in the previous proof, we define $\mathbf{y}'_\ell = \mathbf{y}_\ell \bmod qM$. Then, we have $\mathbf{y}_\ell = \mathbf{y}''_\ell + \mathbf{y}'_\ell$ with \mathbf{y}''_ℓ statistically independent from the z_ℓ 's and distributed as $D_{qM, s, -\mathbf{y}'_\ell}$. By Lemma 2.11, for $s \geq \eta_\varepsilon(qM)$ and $t = \omega(\sqrt{\log N})$, we know that:

$$\Pr_{\forall \ell: \mathbf{y}''_\ell \sim D_{qM, s, -\mathbf{y}'_\ell}} \left[\left\| \sum_{\ell=1}^m z_\ell \cdot (\mathbf{y}''_\ell + \mathbf{y}'_\ell) \right\| \geq st\sqrt{N} \cdot \|\mathbf{z}\| \right] \leq N^{-\omega(1)}.$$

So, with probability close to 1, we have $\|\sum_{\ell=1}^m z_\ell \cdot \mathbf{y}_\ell\| \leq st\sqrt{N} \cdot \|\mathbf{z}\|$. As $0 < \|\mathbf{z}\| \leq \beta$, we obtain:

$$\|\mathbf{h}\| = \frac{1}{q} \left\| \sum_{\ell=1}^m z_\ell \cdot \mathbf{y}_\ell \right\| \leq \frac{st\beta\sqrt{N}}{q}.$$

Finally, since $s \leq \frac{q\|\mathbf{S}\|}{2\beta t\sqrt{N}}$, we obtain $\|\mathbf{h}\| \leq \frac{\|\mathbf{S}\|}{2}$. \square

4 Learning with Errors over modules

In this section, we describe a reduction from Mod-GIVP to M-LWE (Learning With Errors over Modules).

4.1 Variants of LWE

The Learning with Errors Problem. We let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ denote the segment $[0, 1)$ with addition modulo 1. Let us recall the following definitions from [36]. For a probability density function χ on \mathbb{T} and a vector $\mathbf{s} \in \mathbb{Z}_q^n$, we let $A_{\mathbf{s}, \chi}$ denote the distribution on $\mathbb{Z}_q^n \times \mathbb{T}$ obtained by choosing a vector $\mathbf{a} \in \mathbb{Z}_q^n$ uniformly at random, choosing $e \in \mathbb{T}$ according to χ , and returning $(\mathbf{a}, \frac{1}{q}\langle \mathbf{a}, \mathbf{s} \rangle + e)$.

Definition 4.1. *The search version of the Learning With Error problem $\text{SLWE}_{q,\chi}$ is as follows: Let $\mathbf{s} \in \mathbb{Z}_q^n$ be secret; Given arbitrarily many samples from $A_{\mathbf{s},\chi}$, the goal is to find \mathbf{s} .*

The decision version of the Learning With Error problem $\text{LWE}_{q,\chi}$ is as follows: Let $\mathbf{s} \in \mathbb{Z}_q^n$ be uniformly random; The goal is to distinguish between arbitrarily many independent samples from $A_{\mathbf{s},\chi}$ 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 $(\mathbf{a}_i, \frac{1}{q}\langle \mathbf{a}_i, \mathbf{s} \rangle + e_i)$ from $A_{\mathbf{s},\chi}$ is m , then we consider the matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ whose rows are the \mathbf{a}_i 's, and we create the vector $\mathbf{e} = (e_1, \dots, e_m)^T$. Then SLWE is as follows:

$$\begin{array}{c} \updownarrow \\ m \\ \left[\begin{array}{c} \mathbf{A} \\ \hline \end{array} \right] \\ \leftarrow n \end{array} , \quad \frac{1}{q} \cdot \left[\begin{array}{c} \mathbf{A} \\ \hline \end{array} \right] \begin{bmatrix} s_1 \\ \vdots \\ s_n \end{bmatrix} + \begin{bmatrix} e_1 \\ \vdots \\ e_m \end{bmatrix} \xrightarrow{\text{find}} \mathbf{s}$$

Theorem 4.2 ([36]). *Let $\varepsilon(n) = n^{-\omega(1)}$, $\alpha \in (0, 1)$ and $q \geq 2$ such that $\alpha q > 2\sqrt{n}$. There exists a quantum reduction from solving $\text{GIVP}_{\sqrt{8n}/\alpha}^{\eta_\varepsilon}$ in polynomial time (in the worst case, with high probability) to solving SLWE_{q,D_α} in polynomial time with non-negligible probability.*

Assume that q is prime, $q \leq \text{poly}(n)$, and that χ is a probability density function on \mathbb{T} . Then there exists a polynomial-time reduction from $\text{SLWE}_{q,\chi}$ to $\text{LWE}_{q,\chi}$.

The following result from [8] allows to remove the assumption that q is prime.

Theorem 4.3 (Adapted from [8, Cor. 3.2 and 3.3]). *Let $\alpha > 0$, $\varepsilon \in (0, 1/2)$, $m \geq n \geq 1$, $p \geq 25$ and $q \in [p, 2p)$. There exists a polynomial time reduction from $\text{LWE}_{q,\alpha}$ to $\text{LWE}_{p,\Psi_{\leq \beta}}$ where*

$$\beta = C\alpha\sqrt{n}\sqrt{\log(n/\varepsilon)\log(nm/\varepsilon)}$$

for some absolute constant C , and that loses at most ε in the advantage.

LWE over rings. The R-LWE problem was introduced by Lyubashevsky et al. in [16]. Let ψ be a distribution on $\mathbb{T}_{R^\vee} = K_{\mathbb{R}}/R^\vee$ and $s \in R_q^\vee$. We let $A_{s,\psi}^{(R)}$ denote the distribution on $R_q \times \mathbb{T}_{R^\vee}$ obtained by choosing $a \in R_q$ uniformly at random and $e \in \mathbb{T}_{R^\vee}$ according to ψ , and returning $(a, (a \cdot s)/q + e)$.

Definition 4.4. *Let $q \geq 2$ and Ψ be a family of distributions on \mathbb{T}_{R^\vee} . The search version of the Ring Learning With Error problem $\text{R-SLWE}_{q,\Psi}$ is as follows: Let $s \in R_q^\vee$ 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 $\text{R-LWE}_{q,\Upsilon}$ is as follows: Let $s \in R_q^\vee$ be uniformly random and ψ be sampled from Υ ; The goal is to distinguish between arbitrarily many independent samples from $A_{s,\psi}^{(R)}$ and the same number of independent samples from $U(R_q, \mathbb{T}_{R^\vee})$.

As for R-SIS, this problem can be interpreted in terms of linear algebra. In the same example setting as in the case of R-SIS (i.e., the parameter ν is set to a power of 2, implying that $R \simeq \mathbb{Z}[x]/(x^n + 1)$), R-SIS is a variant of LWE where the matrix \mathbf{A} is restricted to being block-negacirculant: $\mathbf{A} = [\text{Rot}(a_1) | \dots | \text{Rot}(a_m)]^T$. The two main results from [16] are a reduction from Id-GIVP to R-SLWE and a reduction from the search version R-SLWE to the decision version R-LWE.

Theorem 4.5 ([16, Th. 4.1 and Th. 5.1]). *Let $\varepsilon(n) = n^{-\omega(1)}$, $\alpha \in (0, 1)$ and $q \geq 2$ of known factorization such that $\alpha q > \omega(\sqrt{\log n})$. There exists a quantum reduction from solving $\text{Id-GIVP}_\gamma^{\eta_\varepsilon}$ in polynomial time (in the worst case, with high probability) to solving $\text{R-SLWE}_{q, \Psi_{\leq \alpha}}$ in polynomial time with non-negligible probability with $\gamma = \sqrt{n} \cdot \omega(\sqrt{\log n})/\alpha$.*

Assume that q is prime, $q \leq \text{poly}(n)$, and that $q \equiv 1 \pmod{\nu}$. Then there exists a polynomial time reduction from $\text{R-SLWE}_{q, \Psi_{\leq \alpha}}$ to $\text{R-LWE}_{q, \mathcal{Y}_\alpha}$.

LWE over modules. The M-LWE problem generalizes both LWE and R-LWE, and was recently introduced in [7]. As in Section 3, the variable n and d respectively denote the dimension of R and the rank of the module $M \subseteq R^d$; We still let $N = nd$ denote the dimension of the corresponding module lattice.

Let ψ be some probability distribution on \mathbb{T}_{R^ν} and $\mathbf{s} \in (R_q^\vee)^d$ be a vector. We define $A_{q, \mathbf{s}, \psi}^{(M)}$ as the distribution on $(R_q)^\nu \times \mathbb{T}_{R^\nu}$ obtained by choosing a vector $\mathbf{a} \in (R_q)^\nu$ uniformly at random, and $e \in \mathbb{T}_{R^\nu}$ according to ψ , and returning $(\mathbf{a}, \frac{1}{q}\langle \mathbf{a}, \mathbf{s} \rangle + e)$.

Definition 4.6. *Let $q \geq 2$ and Ψ be a family of distributions on \mathbb{T}_{R^ν} . The search version of the Module Learning With Error problem $\text{M-SLWE}_{q, \Psi}$ is as follows: Let $\mathbf{s} \in (R_q^\vee)^d$ be secret and $\psi \in \Psi$; Given arbitrarily many samples from $A_{q, \mathbf{s}, \psi}^{(M)}$, the goal is to find \mathbf{s} .*

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

As for LWE and R-LWE, the problem M-LWE can be interpreted in terms of linear algebra. Still in the same example setting, it consists in taking the LWE matrix \mathbf{A} of the form:

We now give our two main results concerning the hardness of M-LWE, in the following two theorems.

Theorem 4.7. *Let $\varepsilon(N) = N^{-\omega(1)}$, $\alpha \in (0, 1)$ and $q \geq 2$ of known factorization such that $\alpha q > 2\sqrt{d} \cdot \omega(\sqrt{\log n})$. There is a quantum reduction from solving $\text{Mod-GIVP}_\gamma^{\eta_\varepsilon}$ in polynomial time (in the worst case, with high probability) to solving $\text{M-SLWE}_{q, \Psi_{\leq \alpha}}$ in polynomial time with non-negligible advantage with $\gamma = \sqrt{8Nd} \cdot \omega(\sqrt{\log n})/\alpha$.*

Assume that q is prime, $q \leq \text{poly}(N)$ and that $q \equiv 1 \pmod{\nu}$. Then there exists a polynomial time reduction from $\text{M-SLWE}_{q, \Psi_{\leq \alpha}}$ to $\text{M-LWE}_{q, \mathcal{Y}_\alpha}$.

In the case of a sub-exponential oracle (and with $\varepsilon(N) = 2^{-\Omega(N)}$), the result still holds and the conditions on the parameters become $\alpha q > 2\sqrt{d} \cdot \Omega(\sqrt{n})$ and $\gamma = d \cdot \Omega(n)/\alpha$.

When $n = N$ and $d = 1$, our theorem is identical to Theorem 4.5 [16, Th. 3.1]. When $n = 1$ and $d = N$, it is identical to Theorem 4.2 [36, Th. 4.1 and 5.1] (apart from a minor detail with the error distribution which can easily be handled: we use \mathcal{Y}_α rather than D_α).

Theorem 4.8. *Let $p, q \in [2, 2^{N^{O(1)}}]$ and $\alpha, \beta \in (0, 1)$ such that $\beta \geq \alpha \cdot \max(1, \frac{q}{p}) \cdot n^{1/4} N^{1/2} \cdot \omega(\log^2 N)$ and $\alpha q \geq \omega(\sqrt{\log(N)/n})$. There exists a polynomial time reduction from M-LWE $_{q, r_\alpha}$ to M-LWE $_{p, r_\beta}$.*

In the case of a sub-exponential oracle (and with $\varepsilon(N) = 2^{-\Omega(N)}$), the result still holds and the conditions on the parameters become $\beta \geq \alpha \cdot \max(1, \frac{q}{p}) \cdot \Omega(n^{1/4} N^{5/2})$ and $\alpha q \geq \Omega(\sqrt{d})$.

Note that the condition on αq from Theorem 4.8 is always weaker than the one from Theorem 4.7. Combined with Theorem 4.7 by using a q prime close to p with $q = 1 \pmod{\nu}$, Theorem 4.8 provides a reduction from Mod-SIVP to M-LWE with a modulus p of arbitrary arithmetic form. As M-LWE is a generalization of both LWE and R-LWE, this theorem also provides a reduction from Id-SIVP $_\gamma$ to RLWE $_{p, r_\beta}$, for a modulus p of arbitrary arithmetic shape. Note that in the case of LWE, this theorem is almost identical to Theorem 4.3.

The remainder of this section is devoted to proving Theorems 4.7 (Subsection 4.2 and 4.3) and 4.8 (Subsection 4.4).

4.2 A reduction from Mod-GIVP to M-SLWE.

The reduction from Mod-GIVP to M-SLWE follows the same design principle as Regev's reduction from GIVP to SLWE. It makes use of the following intermediate problems, where ϕ denotes an arbitrary real-valued function on lattices and where γ is a function of the dimension, called *Module Discrete Gaussian Sampling problem* (M-DGS $_\gamma^\phi$): Given an N -dimensional module lattice M and a number $r > \gamma \cdot \phi(M)$, the goal is to output a sample from $D_{M, r}$. The reduction proceeds in two steps:

$$\text{Mod-GIVP}_{\frac{\eta_\varepsilon}{\sqrt{8Nd \cdot \omega(\sqrt{\log n})/\alpha}}} \xrightarrow{\text{Lemma 4.9}} \text{M-DGS}_{\frac{\eta_\varepsilon}{\sqrt{2d \cdot \omega(\sqrt{\log n})/\alpha}}} \xrightarrow{\text{Lemma 4.10}} \text{M-SLWE}_{q, \psi_{\leq \alpha}}$$

The first reduction comes directly from the reduction from GIVP to DGS given by [36, Le. 3.17]: It is lattice-preserving and thus also works when we consider the problems restricted to any family of lattices.

Lemma 4.9 (Adapted from [36, Le. 3.17]). *For any $\varepsilon = \varepsilon(N) \leq 1/10$ and any γ and ϕ such that $\gamma \cdot \phi(M) \leq \sqrt{2}\eta_\varepsilon(M)$, there is a polynomial time reduction from Mod-GIVP $_{\frac{\phi}{2\sqrt{N} \cdot \gamma}}$ to M-DGS $_\gamma^\phi$.*

In contrast, the second one needs to be adapted.

Lemma 4.10. *Let $\varepsilon(N) = N^{-\omega(1)}$, $\alpha \in (0, 1)$ and q be some integer such that $\alpha q \geq 2\sqrt{d} \cdot \omega(\sqrt{\log n})$. Assume that we have access to an oracle that solves M-SLWE $_{q, \psi_{\leq \alpha}}$ given a polynomial number of samples. Then there exists a polynomial time quantum algorithm for M-DGS $_{\frac{\eta_\varepsilon}{\sqrt{2d \cdot \omega(\sqrt{\log n})/\alpha}}}$.*

Proof. We use the same principle as Regev's reduction [36, Th. 3.1]. We consider a module lattice M and a number $r \geq \sqrt{2d} \cdot \omega(\sqrt{\log n}) \cdot \eta_\varepsilon(M)/\alpha$. The idea is to produce samples for $D_{M, r'}$ with r' large enough, and then to use Lemma 4.11 several times to progressively decrease the value of r' . Take $r_i = r \cdot (\alpha q / \sqrt{d} \cdot \omega(\sqrt{\log n}))^i$. The first iteration starts with $r_{3N} > 2^{3N} > 2^{2N} \lambda_N(M)$ (using a LLL-reduction algorithm beforehand). Then it obtains $\text{poly}(N)$ samples of $D_{M, r_{3N}}$ using the algorithm of Theorem 2.3, and finishes with $\text{poly}(N)$ samples of $D_{M, r_{3N-1}}$ (the reduction repeats $\text{poly}(N)$ times the same iteration with the same samples in input to obtain sufficiently many different samples in output). It iterates until having $\text{poly}(N)$ samples of D_{M, r_1} with $r_1 = r \cdot \alpha q / (\sqrt{d} \cdot \omega(\sqrt{\log n})) > \sqrt{2}q \cdot \eta_\varepsilon(M)$ then it iterates a last time to obtain samples of D_{M, r_0} with $r_0 = r > \sqrt{d} \cdot \omega(\sqrt{\log n}) \cdot \eta_\varepsilon(M)/\alpha$. These samples are solutions to M-DGS $_{\frac{\eta_\varepsilon}{\sqrt{2d \cdot \omega(\sqrt{\log n})/\alpha}}}$. \square

We now describe the iterative step:

Lemma 4.11. *Let $\varepsilon(N) = N^{-\omega(1)}$, $\alpha \in (0, 1)$ and $q \geq 2$. Assume that we have access to an oracle that solves M-SLWE $_{q, \psi_{\leq \alpha}}$ in polynomial time with non-negligible probability. Then there exists a polynomial time quantum algorithm that, given an N -dimensional module lattice M , a number $r > \sqrt{2}q \cdot \eta_\varepsilon(M)$ and $\text{poly}(N)$ samples from $D_{M, r}$, produces a sample from $D_{M, \frac{r \cdot \sqrt{d} \cdot \omega(\sqrt{\log n})}{\alpha q}}$ with non-negligible probability.*

To prove Lemma 4.11, we use the intermediate problem Mod-BDD_δ : Given a module lattice M , $\delta < \lambda_1(M)/2$ and any point $\mathbf{y} \in \mathbb{R}^n$ of the form $\mathbf{y} = \mathbf{x} + \mathbf{e}$ for some $\mathbf{x} \in M$ and $\|\mathbf{e}\|_{2,\infty} \leq \delta$, find \mathbf{x} . Note that we use the $\ell_{2,\infty}$ rather than the euclidean norm, as it is more convenient in Lemma 4.14.

As in [36], we use another intermediate problem called $q\text{-Mod-BDD}_\delta$: Given a module lattice M and a point $\mathbf{y} \in \mathbb{R}^n$ within distance (with respect to $\ell_{2,\infty}$ norm) δ of M , output the coset in M/qM of the closest vector to \mathbf{y} . The proof of Lemma 4.11 consists of a sequence of reductions (note that δ is set to $\frac{\alpha q \cdot \omega(\sqrt{\log n})}{\sqrt{2nr}}$).

$$\text{Samples from } D_{M, \frac{r\sqrt{d}\cdot\omega(\sqrt{\log n})}{\alpha q}} \xrightarrow[\text{(quantum)}]{\text{Lemma 4.12}} \text{Mod-BDD}_{M^\vee, \delta} \xrightarrow{\text{Lemma 4.13}} q\text{-Mod-BDD}_{M^\vee, \delta} \xrightarrow{\text{Lemma 4.14}} \text{M-SLWE}_{q, D_\alpha} + \text{Samples from } D_{M,r}$$

The first reduction of Regev's proof is quantum and also lattice-preserving. It is adapted to the $\ell_{2,\infty}$ norm rather than the euclidean norm (note that $\lambda_1(M^\vee)$ is still with respect to the euclidean norm). For the adaptation, we used the fact that an N -dimensional vector sampled from D_s has $\ell_{2,\infty}$ norm at most $s\sqrt{d}\omega(\sqrt{\log n})$, except with negligible probability.

Lemma 4.12 (Adapted from [36, Le. 3.14]). *There exists an efficient quantum algorithm that, given any N -dimensional module lattice M , a number $\delta < \lambda_1(M^\vee)\omega(\sqrt{\log n})/(2\sqrt{n})$, and an oracle that solves Mod-BDD_δ on M^\vee , outputs samples from $D_{M, \sqrt{d}\omega(\sqrt{\log n})/(\sqrt{2}\delta)}$.*

Note that by Lemma 2.5, as $r > \sqrt{2}q \cdot \eta_\varepsilon(M)$, we have that $\delta = \frac{\alpha q \cdot \omega(\sqrt{\log n})}{\sqrt{2nr}} < \frac{\omega(\sqrt{\log n})}{\sqrt{n} \cdot \eta_\varepsilon(M)} < \frac{\lambda_1(M^\vee)\omega(\sqrt{\log n})}{2\sqrt{n}}$.

The second reduction is a special case of [36, Le. 3.5], which is lattice-preserving (and hence also applies to module lattices).

Lemma 4.13 ([36, Le. 3.5]). *For any $q \geq 2$, there is a polynomial time reduction from Mod-BDD_δ to $q\text{-Mod-BDD}_\delta$.*

We will modify the last reduction, by proving the following adaptation of [36, Le. 3.4]. The following lemma is the main modification of the proof of the first part of Theorem 4.7.

Lemma 4.14. *Let $\varepsilon(N) = N^{-\omega(1)}$, $\alpha \in (0, 1)$ 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 $q\text{-Mod-BDD}_{M^\vee, \frac{\alpha q \cdot \omega(\sqrt{\log n})}{\sqrt{2nr}}}$ to $\text{M-SLWE}_{q, \Psi_{\leq \alpha}}$.*

The principle of the reduction is to construct from \mathbf{y} , the input of $q\text{-Mod-BDD}$, and from some discrete and continuous Gaussian samples, the pairs (\mathbf{a}, b) distributed as $A_{q, \mathbf{s}, \psi}^{(M)}$, where \mathbf{s} will directly depend on the closest vector \mathbf{x} to \mathbf{y} . To produce such samples (\mathbf{a}, b) with the desired distribution, we combine the corresponding proofs for LWE and R-LWE (those of Lemmata [36, Le. 3.4] and [16, Le. 4.5]). Then a call to the oracle of M-SLWE returns \mathbf{s} and lets us recover information on \mathbf{x} .

Proof of Lemma 4.14. Let \mathcal{O} be the oracle which, given $m \leq \text{poly}(N)$ samples (\mathbf{a}, b) from $A_{q, \mathbf{s}, \psi}^{(M)}$ for $\psi \in \Psi_{\leq \alpha}$, outputs \mathbf{s} in polynomial time with probability $N^{-O(1)}$. Given $M = \sum_{k=1}^d I_k \cdot \mathbf{b}_k$, the input of the reduction is $\mathbf{y} = \mathbf{x} + \mathbf{e}$ such that $\mathbf{x} \in M^\vee$ and $\|\mathbf{e}\|_{2,\infty} \leq \delta = \frac{\alpha q \cdot \omega(\sqrt{\log n})}{\sqrt{2nr}}$. The goal is to find $\mathbf{x} \bmod qM^\vee$. The reduction is as follows:

- For all $\ell \leq m$:
 - Get a fresh \mathbf{z}_ℓ distributed as $D_{M,r}$ and a fresh e'_ℓ distributed as $D_{\alpha/\sqrt{2}}$,
 - Let $\mathbf{a}_\ell = \Theta^{-1}(\mathbf{z}_\ell \bmod qM)$ and $b_\ell = \frac{1}{q}\langle \mathbf{z}_\ell, \mathbf{y} \rangle + e'_\ell \bmod R^\vee$ (see the definition of Θ in Section 2.1).
- Invoke the oracle \mathcal{O} on input $\{(\mathbf{a}_\ell, b_\ell)\}_{\ell \leq m}$. If \mathcal{O} succeeds, it returns some $\mathbf{s} \in (R_q^\vee)^d$.

- Output $\Theta^{-1}(\mathbf{s}) \in M^\vee/qM^\vee$.

We show that the oracle \mathcal{O} is used properly, i.e., that its input follows the distribution $A_{q,\mathbf{s},\psi}^{(M)}$.

Lemma 4.15. *Let $\varepsilon > 0$ and $\mathbf{s} = \Theta(\mathbf{x} \bmod qM^\vee)$. There exists $\psi \in \Psi_{\leq \alpha}$ such that the statistical distance between $A_{q,\mathbf{s},\psi}^{(M)}$ and the distribution of (\mathbf{a}, b) is at most 6ε .*

Proof. We first show that the statistical distance between \mathbf{a} , the first component of each sample, and the uniform distribution on R_q^d is at most 2ε . By Lemma 2.6, the statistical distance between the distribution of \mathbf{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 $\mathbf{a} = \Theta^{-1}(\mathbf{z} \bmod qM)$ and the uniform distribution on $(R_q)^d$ is at most 2ε .

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

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

By definition, we have $\mathbf{z} = \Theta(\mathbf{a}) = \sum_{k=1}^d (t_k \cdot a_k) \cdot \mathbf{b}_k \bmod qM$ with $t_k \in I_k$ and $a_k \in R_q$. By Lemma 2.1, we have $M^\vee = \sum_{k=1}^d I_k^\vee \cdot \mathbf{b}_k^\vee$. Let $\mathbf{x} = \sum_{k=1}^d x_k \cdot \mathbf{b}_k^\vee$. We have that $x_k \in I_k^\vee = I_k^{-1} \cdot R^\vee$ for all k . We also have $\langle \mathbf{b}_k, \mathbf{b}_{k'}^\vee \rangle = 1$ if $k = k'$ and $\langle \mathbf{b}_k, \mathbf{b}_{k'}^\vee \rangle = 0$ otherwise. Then, modulo qR^\vee :

$$\langle \mathbf{z}, \mathbf{x} \rangle = \sum_{k,k'=1}^d (t_k \cdot a_k) \cdot x_{k'} \cdot \langle \mathbf{b}_k, \mathbf{b}_{k'}^\vee \rangle = \sum_{k=1}^d (t_k \cdot a_k) \cdot x_k = \sum_{k=1}^d a_k \cdot (t_k \cdot x_k).$$

Because $\mathbf{s} = \Theta(\mathbf{x} \bmod qM^\vee) = (t_1 \cdot x_1 \bmod qR^\vee, \dots, t_d \cdot x_d \bmod qR^\vee)^T$, we have:

$$\langle \mathbf{a}, \mathbf{s} \rangle = \sum_{k=1}^d a_k \cdot (t_k \cdot x_k) = \langle \mathbf{z}, \mathbf{x} \rangle \bmod qR^\vee.$$

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

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

We then show that e' is following the same distribution as $\langle \mathbf{e}'', \mathbf{e} \rangle$ with $\mathbf{e}'' \sim D_{\mathbf{s}, \dots, \mathbf{s}}$, $\mathbf{s} = (s_i)_i$ and $s_i = s_{\nu-i} = \alpha/\sqrt{2\sum_{k=1}^d |\sigma_i(e_k)|^2}$ for $i \in \mathbb{J}$. By Lemma 2.13, as the vector \mathbf{e}'' is distributed from $D_{\mathbf{s}, \dots, \mathbf{s}}$ and $\mathbf{e} \in K^d$ is fixed, we have that $\langle \mathbf{e}'', \mathbf{e} \rangle$ has distribution $D_{\mathbf{s}'}$ with $s'_i = s'_{\nu-i} = s_i \sqrt{\sum_{k=1}^d |\sigma_i(e_k)|^2} = \alpha/\sqrt{2}$.

We are now led to considering the distribution of $\langle \mathbf{z}' + \mathbf{e}'', \mathbf{e} \rangle$. We write $\mathbf{e}'' = \mathbf{e}''_1 + \mathbf{e}''_2$ with $\mathbf{e}''_1 \sim D_{\alpha/(\sqrt{2}\delta)}$ and $\mathbf{e}''_2 \sim D_{\mathbf{s}''}$ with $(s''_i)^2 = s_i^2 - \alpha^2/(2\delta^2)$ (which is positive, by the assumption on $\|\mathbf{e}\|_{2,\infty}$). As we have $\alpha/(\sqrt{2}\delta) = r/q$ and $r/q \geq \sqrt{2}\eta_\varepsilon(M)$, Lemma 2.8 gives us that the statistical distance between the distribution of $\mathbf{z}' + \mathbf{e}''_1$ and $D_{\alpha/\delta}$ is at most 4ε . As a consequence, the statistical distance between the distribution of $\mathbf{z}' + \mathbf{e}''_1 + \mathbf{e}''_2$ and $D_{r'', \dots, r''}$ 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 2.13 again with the fixed vector \mathbf{e} , we obtain that the statistical distance between the distribution of $\langle \mathbf{z} + \mathbf{e}'', \mathbf{e} \rangle$ and $D_{r'}$ 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. \square

As the input of \mathcal{O} is within negligible statistical distance from $A_{q,\mathbf{s},\psi}^{(M)}$ for a distribution $\psi \in \Psi_{\leq \alpha}$ and $\mathbf{s} = \Theta(\mathbf{x} \bmod qM^\vee)$, oracle \mathcal{O} succeeds with non-negligible probability. If it does succeed, then the output of our reduction is $\mathbf{x} \bmod qM^\vee$, which completes the proof of Lemma 4.14. \square

This concludes the proof of the first part of Theorem 4.7.

4.3 A reduction from M-SLWE to M-LWE.

We now describe a reduction from the search version M-SLWE to the decision version M-LWE. The reduction of Regev from SLWE to LWE in [36] does not carry over to the structured variants of LWE. We instead use the line of proof of Lyubashevsky et al. in [16]. Let $q = 1 \bmod \nu$ be prime. Then $(q) = \prod_{i \in \mathbb{Z}_\nu^\times} \mathfrak{q}_i$ where any \mathfrak{q}_i is a prime ideal with norm $\mathcal{N}(\mathfrak{q}_i) = q$. For $i \in \mathbb{Z}_\nu^\times$, we let i^- denote the largest element in \mathbb{Z}_ν^\times less than i (and we define 1^- as 0). We define the following intermediate problems:

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

We consider the following sequence of reductions:

$$\text{M-SLWE}_{q,\Psi} \xrightarrow{\text{Lemma 4.16}} \mathfrak{q}_i\text{-MLWE}_{q,\Psi} \xrightarrow{\text{Lemma 4.18}} \text{DecMLWE}_{q,\Psi}^i \xrightarrow{\text{Lemma 4.19}} \text{M-DLWE}_{q,\mathcal{Y}}^i \xrightarrow{\text{Lemma 4.21}} \text{M-LWE}_{q,\mathcal{Y}}$$

We explain the first two reductions. The following result is adapted from [16].

Lemma 4.16. *For any $i \in \mathbb{Z}_\nu^\times$, there is a polynomial time reduction from $\text{M-SLWE}_{q,\Psi_{\leq \alpha}}$ to $\mathfrak{q}_i\text{-MLWE}_{q,\Psi_{\leq \alpha}}$.*

Proof. We aim at using an oracle solving $\mathfrak{q}_i\text{-MLWE}$ for finding the values of $\mathbf{s} \bmod \mathfrak{q}_j R^\vee$ for every $j \in \mathbb{Z}_\nu^\times$. Then, by the Chinese Remainder Theorem, this allows us to construct $\mathbf{s} \bmod R^\vee$ and to solve M-SLWE.

We use the K -automorphisms, defined by $\tau_j(\xi) = \xi^j$ for all $j \in \mathbb{Z}_\nu^\times$. We choose $j_i \in \mathbb{Z}_\nu^\times$ such that $\tau_{j_i}(\mathfrak{q}_j) = \mathfrak{q}_i$. The reduction is as follows:

- For every sample (\mathbf{a}, \mathbf{b}) , create the sample $(\mathbf{a}', \mathbf{b}')$ with $\mathbf{a}' = (\tau_{j_i}(a_1), \dots, \tau_{j_i}(a_d))^T$ and $\mathbf{b}' = \tau_{j_i}(\mathbf{b})$.
- Use the oracle of $\mathfrak{q}_i\text{-MLWE}$ with these samples, and get $\mathbf{t} \in (R^\vee/\mathfrak{q}_i R^\vee)^d$.
- Return $(\tau_{j_i}^{-1}(t_1), \dots, \tau_{j_i}^{-1}(t_d)) \in (R^\vee/\mathfrak{q}_j R^\vee)^d$.

We show that $\tau_{j_i}^{-1}(t_k) = s_k \bmod \mathfrak{q}_j R^\vee$ for all $k \in [d]$. By definition, we have $\mathbf{b} = \frac{1}{q} \langle \mathbf{a}, \mathbf{s} \rangle + e \bmod R^\vee$ with $\langle \mathbf{a}, \mathbf{s} \rangle = \sum_{k=1}^d a_k \cdot s_k$. As a consequence, we have:

$$\mathbf{b}' = \tau_{j_i}(\mathbf{b}) = \frac{1}{q} \sum_{k=1}^d \tau_{j_i}(a_k) \cdot \tau_{j_i}(s_k) + \tau_{j_i}(e) = \frac{1}{q} \langle \mathbf{a}', \mathbf{s}' \rangle + \tau_{j_i}(e) \bmod R^\vee,$$

with $\mathbf{s}' = (\tau_{j_i}(s_1), \dots, \tau_{j_i}(s_d))^T$. As τ_{j_i} is an automorphism, the vector \mathbf{a}' is uniformly distributed in R_q^d . Also, as $\Psi_{\leq \alpha}$ is closed under the automorphisms of K (see Lemma 4.17), we have $\psi' := \tau_{j_i}(\psi) \in \Psi_{\leq \alpha}$.

Overall, the pairs (\mathbf{a}', b') are distributed as $A_{q, \mathbf{s}', \psi'}^{(M)}$. If successful, the \mathfrak{q}_i -MLWE oracle outputs $\mathbf{t} = \mathbf{s}' \bmod \mathfrak{q}_i R^\vee = (\tau_{j_i}(s_1) \bmod \mathfrak{q}_i R^\vee, \dots, \tau_{j_i}(s_d) \bmod \mathfrak{q}_i R^\vee)$. Then our reduction returns $(\tau_{j_i}^{-1}(t_1), \dots, \tau_{j_i}^{-1}(t_d))^T \in (R^\vee / \mathfrak{q}_j R^\vee)^d$, which is equal to $\mathbf{s} \bmod \mathfrak{q}_j R^\vee$. \square

By [16, Le. 5.6], we know that $\Psi_{\leq \alpha}$ satisfies the property required by Lemma 4.16.

Lemma 4.17 ([16, Le. 5.6]). *For any $\alpha > 0$, the family $\Psi_{\leq \alpha}$ is closed under every automorphism τ of K , i.e., $\psi \in \Psi_{\leq \alpha} \Rightarrow \tau(\psi) \in \Psi_{\leq \alpha}$.*

We now describe the next reduction.

Lemma 4.18. *For any $i \in \mathbb{Z}_\nu^\times$, there is a polynomial time reduction from \mathfrak{q}_i -MLWE $_{q, \Psi}$ to DecMLWE $_{q, \Psi}^i$.*

Proof. We want to find $\mathbf{s} \bmod \mathfrak{q}_i R^\vee$ from samples from $A_{q, \mathbf{s}, \psi}^{(M)}$, by using an oracle that solves the DecMLWE $_{q, \Psi}^i$ problem. The principle of the proof is to find, one by one, each one of the d coordinates of $\mathbf{s} \bmod \mathfrak{q}_i R^\vee$ by using the oracle of DecMLWE $_{q, \Psi}^i$. For each coordinate, there are $\mathcal{N}(\mathfrak{q}_i) = q \leq \text{poly}(n)$ 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 [36, Le. 4.2] and randomize a coordinate of \mathbf{a} .

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

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

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

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

$$b' = b + \frac{1}{q}(r + xy) = \frac{1}{q} \left(\sum_{k=1}^d a_k \cdot s_k + r + xy \right) + e = \left(\frac{1}{q}(\mathbf{a}', \mathbf{s}) + e \right) + \frac{1}{q}(r + y(x - s_1)).$$

We have two cases:

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

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

The last reductions carry over directly from the ring setting [16, Le. 5.12 and 5.14] to the module setting (the proof randomizes the noise distribution Ψ , which is the same in the ring and module settings).

Lemma 4.19 (Adapted from [16, Le. 5.12]). *For any $\alpha > 0$ and every $i \in \mathbb{Z}_\nu^\times$, there is a polynomial time reduction from Dec-MLWE $_{q, \Psi_{\leq \alpha}}^i$ to M-DLWE $_{q, \Upsilon_\alpha}^i$.*

Lemma 4.20 (Adapted from [16, Le. 5.13]). *Let $\alpha > (1/q)\eta_\varepsilon(R^\vee)^d$ for some ε . Then for any ψ in the support of Υ_α and $\mathbf{s} \in (R^\vee)^d$, the distribution $A_{q, \mathbf{s}, \psi}^{(M, \nu-1)}$ is within statistical distance $\varepsilon/2$ of the uniform distribution over $((R_q)^\vee, \mathbb{T}_{R^\vee})$.*

Lemma 4.21 (Adapted from [16, Le. 5.14]). *Let \mathcal{Y} be a distribution over noise distributions satisfying that for any ψ in the support of \mathcal{Y} and any $\mathbf{s} \in (R_q^\vee)^d$, the distribution $A_{q,\mathbf{s},\psi}^{(M,\nu-1)}$ is within negligible statistical distance from uniform. Then for any oracle solving the M-LWE $_{q,\mathcal{Y}}$ problem, there exists an $i \in \mathbb{Z}_\nu^\times$ and an efficient algorithm that solves the M-DLWE $_{q,\mathcal{Y}}^i$ using the oracle.*

This completes the proof of Theorem 4.7.

4.4 A modulus-switching self-reduction for M-LWE

The aim of the present section is to give the proof of Theorem 4.8: For any $p, q \geq 2$, and under some conditions on α and β , M-LWE $_{p,\mathcal{Y}_\beta}$ is no easier than M-LWE $_{q,\mathcal{Y}_\alpha}$. We proceed by a sequence of reductions:

$$\text{M-LWE}_{q,\mathcal{Y}_\alpha} \xrightarrow[4.22 \text{ and } 4.24]{\text{Lemmata}} \text{HNF-MLWE}_{q,D_{\frac{1}{q}R^\vee, [\alpha, \alpha']}} \xrightarrow{\text{Lemma 4.26}} \text{M-LWE}_{p,\Psi_{\leq \beta}} \xrightarrow{\text{Lemma 4.29}} \text{M-LWE}_{p,\mathcal{Y}_\beta}$$

In Lemmata 4.22 and 4.24, we first reduce M-LWE $_{q,\mathcal{Y}_\alpha}$ to the HNF version of M-LWE $_{q,D_{(1/q)R^\vee, [\alpha, \alpha']}}$ (i.e., with a secret \mathbf{s} of small euclidean norm), where $\alpha' \approx \alpha n^{1/4}$. Then, in Lemma 4.26 we reduce HNF-MLWE $_{q,D_{(1/q)R^\vee, [\alpha, \alpha']}}$ to M-LWE $_{p,\Psi_{\leq \beta}}$, by switching the modulus and handling the right hand sides of the M-LWE samples so that the distribution of the error term belongs to $\Psi_{\leq \beta}$. Finally, in Lemma 4.29, we re-randomize the noise distribution, thus providing a reduction from M-LWE $_{p,\Psi_{\leq \beta}}$ to M-LWE $_{p,\mathcal{Y}_\beta}$.

Reducing M-LWE $_{q,\mathcal{Y}_\alpha}$ to HNF-MLWE $_{q,D_{(1/q)R^\vee, [\alpha, \alpha']}}$. We first reduce M-LWE $_{q,\mathcal{Y}_\alpha}$ to M-LWE $_{q,\Psi_{[\alpha, \alpha']}}$. We consider a sample ϕ from \mathcal{Y}_α : we have $\phi = D_{\mathbf{r}}$ with $r_i = r_{\nu-i} = \alpha\sqrt{1 + \sqrt{n}x_i}$ and x_i sampled from $\Gamma(2, 1)$, for all $i \in \mathbb{J}$. By definition of $\Gamma(2, 1)$, we have $\Pr_{x \leftarrow \Gamma(2, 1)}[x \leq t] = 1 - (1+t)e^{-t}$, from which we derive that $x \leq \omega(\log N)$ with probability negligibly close to 1. As a consequence, with the same probability we have that $\alpha < r_i \leq \alpha' = \alpha \cdot n^{1/4} \omega(\log N)$ for all i . Therefore, M-LWE $_{q,\Psi_{[\alpha, \alpha']}}$ is no easier than M-LWE $_{q,\mathcal{Y}_\alpha}$.

Now, for any distribution $D_{\mathbf{r}}$ arbitrarily chosen in $\Psi_{[\alpha, \alpha]}$, we discretize the noise distribution by proving that M-LWE $_{q,D_{(1/q)R^\vee, \sqrt{2}\mathbf{r}}}$ is no easier than M-LWE $_{q,D_{\mathbf{r}}}$. Here, by abuse of notation, M-LWE $_{q,D_{(1/q)R^\vee, \sqrt{2}\mathbf{r}}}$ denotes the M-LWE problem where the distribution $\psi = D_{(1/q)R^\vee, \sqrt{2}\mathbf{r}}$ is a discrete distribution on $(1/q)R^\vee$ and where the goal is to distinguish between arbitrarily many independent samples from $A_{q,\mathbf{s},\psi}^{(M)}$ and the same number of independent samples from $U(R_q^d \times \mathbb{T}_{q,R^\vee})$, with $\mathbb{T}_{q,R^\vee} = ((1/q)R^\vee)/R^\vee$.

Lemma 4.22 (Adapted from [12, Le. 2]). *For any $q \geq 2$, $\varepsilon \in (0, 1)$, $\mathbf{r} \in (\mathbb{R}^+)^n$ with $r_{\nu-i} = r_i$ for all i , and $\alpha \in [\eta_\varepsilon(R^\vee)/q, \min_i r_i]$, there is a polynomial time reduction from M-LWE $_{q,D_{\mathbf{r}}}$ to M-LWE $_{q,D_{(1/q)R^\vee, \sqrt{2}\mathbf{r}}}$.*

The proof is following the same design as the proof of [12, Le. 2].

Proof. We consider the following transformation: Given $(\mathbf{a}, b) \in R_q^d \times \mathbb{T}_{R^\vee}$, sample $f \leftarrow D_{(1/q)R^\vee - b, \mathbf{r}}$ and returns $(\mathbf{a}, b + f \bmod R^\vee)$.

If the sample (\mathbf{a}, b) is uniform over $R_q^d \times \mathbb{T}_{R^\vee}$, then $(b + f \bmod R^\vee)$ is uniform in \mathbb{T}_{q,R^\vee} . Now, assume that (\mathbf{a}, b) is distributed according to $A_{q,\mathbf{s},D_{\mathbf{r}}}^{(M)}$: We have $b = \frac{1}{q}\langle \mathbf{a}, \mathbf{s} \rangle + e$, where $e \sim D_{\mathbf{r}}$. Since $\frac{1}{q}\langle \mathbf{a}, \mathbf{s} \rangle$ belongs to $\frac{1}{q}R^\vee$, we have $D_{(1/q)R^\vee - b, \mathbf{r}} = D_{(1/q)R^\vee - e, \mathbf{r}}$. By [30, Th. 3.1], as e is sampled from $D_{\mathbf{r}}$ and $e' = e + f$ with f sampled from $D_{(1/q)R^\vee - e, \mathbf{r}}$, the distribution of e' is statistically close to $D_{(1/q)R^\vee, \sqrt{2}\mathbf{r}}$. We conclude that, in this case, the transformation returns a sample of $A_{q,\mathbf{s},D_{(1/q)R^\vee, \sqrt{2}\mathbf{r}}}^{(M)}$. \square

Finally, Lemma 4.24 allows us to reduce the M-LWE $_{q,D_{(1/q)R^\vee, \sqrt{2}\mathbf{r}}}$ problem to a variant in which the secret is chosen from $D_{(R^\vee)^d, \sqrt{2}q\mathbf{r}}$. We call this new problem the Hermite Normal Form (HNF) of M-LWE.

Definition 4.23. *Let $q \geq 2$, and \mathcal{Y} be a set of distributions over $(1/q)R^\vee$. The Hermite Normal Form of the decision version of the Module Learning With Error problem HNF-MLWE $_{q,\mathcal{Y}}$ is as follows: Let ψ be arbitrarily chosen from \mathcal{Y} and $\mathbf{s} \in (R_q^\vee)^d$ be sampled from $(q \cdot \psi)^d$. The goal is to distinguish between arbitrarily many independent samples from $A_{q,\mathbf{s},\psi}^{(M)}$ and the same number of independent samples from $U(R_q^d \times \mathbb{T}_{q,R^\vee})$.*

We have the following result:

Lemma 4.24 (Adapted from [3, Le. 2]). *There is a polynomial time transformation that, for arbitrary $\mathbf{s} \in (R_q^\vee)^d$ and error distribution $D_{(1/q)R^\vee, \mathbf{r}}$, maps $A_{q, \mathbf{s}, D_{(1/q)R^\vee, \mathbf{r}}}^{(M)}$ to $A_{q, \bar{\mathbf{x}}, D_{(1/q)R^\vee, \mathbf{r}}}^{(M)}$ with $\bar{\mathbf{x}} \leftarrow D_{(R^\vee)^d, q\mathbf{r}}$, and maps $U(R_q^d \times \mathbb{T}_{q, R^\vee})$ to itself.*

The proof is following the same principle as the proof of [3, Le. 2].

Proof. We are given samples from a distribution D that is either the uniform over $R_q^d \times \mathbb{T}_{q, R^\vee}$, or $A_{q, \mathbf{s}, D_{(1/q)R^\vee, \mathbf{r}}}^{(M)}$.

In a first stage, we take several samples (\mathbf{a}, b) from D and construct a set of d pairs $\{(\mathbf{a}_k, b_k)\}$ such that the \mathbf{a}_k 's are linearly independent over R_q and generate R_q^d (recall that R_q is not a field). A polynomial number of samples suffices to obtain such \mathbf{a}_k 's. This can be observed by considering the CRT components of $R_q \simeq (\mathbb{F}_{q^\ell})^{n/\ell}$ independently: An equivalent condition is that the n/ℓ matrices corresponding to each component are invertible over the corresponding finite field. We define $\bar{\mathbf{A}} = (\mathbf{a}_1^T, \dots, \mathbf{a}_d^T)$ and $\bar{\mathbf{b}} = (b_1, \dots, b_d)^T$. By construction, the map $\mathbf{y} \mapsto \bar{\mathbf{A}}\mathbf{y}$ is a bijection of R_q^d , and if $D = A_{q, \mathbf{s}, D_{(1/q)R^\vee, \mathbf{r}}}^{(M)}$ then we have $\bar{\mathbf{b}} = \frac{1}{q}(\bar{\mathbf{A}}\mathbf{s} + \bar{\mathbf{x}})$, where $\bar{\mathbf{x}}$ is sampled from $D_{(R^\vee)^d, q\mathbf{r}}$.

In a second stage, we map the fresh samples (\mathbf{a}, b) from D , to samples (\mathbf{a}', b') with $\mathbf{a}' = -(\bar{\mathbf{A}})^{-T} \cdot \mathbf{a} \in R_q^d$ and $b' = b + \langle \mathbf{a}', \bar{\mathbf{b}} \rangle$. As the map $\mathbf{y} \mapsto \bar{\mathbf{A}}\mathbf{y}$ is a bijection of R_q^d and as \mathbf{a} is uniform in R_q^d , we have that \mathbf{a}' is uniform in R_q^d . For the right hand side b' , we consider two cases:

- If D is the uniform distribution on $R_q^d \times \mathbb{T}_{q, R^\vee}$, then (\mathbf{a}', b') is also uniform on $(R_q)^d \times \mathbb{T}_{q, R^\vee}$.
- If D is $A_{q, \mathbf{s}, D_{(1/q)R^\vee, \mathbf{r}}}^{(M)}$, then $b' = \frac{1}{q}\langle \mathbf{a}', \mathbf{s} \rangle + e - \frac{1}{q}\langle (\bar{\mathbf{A}})^{-T} \mathbf{a}, \bar{\mathbf{A}}\mathbf{s} \rangle + \frac{1}{q}\langle \mathbf{a}', \bar{\mathbf{x}} \rangle = \frac{1}{q}\langle \mathbf{a}', \bar{\mathbf{x}} \rangle + e$. As a consequence, the pair (\mathbf{a}', b') is distributed as $A_{\bar{\mathbf{x}}, D_{(1/q)R^\vee, \mathbf{r}}}^{(q)}$, with $\bar{\mathbf{x}}$ sampled from $D_{(R^\vee)^d, q\mathbf{r}}$. □

This completes the reduction from M-LWE $_{q, \mathcal{Y}_\alpha}$ to HNF-MLWE $_{q, D_{(1/q)R^\vee, [\alpha, \alpha']}}$.

Reducing HNF-MLWE $_{q, D_{(1/q)R^\vee, [\alpha, \alpha']}}$ to M-LWE $_{p, \Psi_{\leq \beta}}$. This is the main component of the proof of Theorem 4.8. In Lemma 4.25, we first give a bound on $\|\mathbf{s}\|_{2, \infty}$.

Lemma 4.25. *Let $\varepsilon = N^{-\omega(1)}$, $\alpha' > \alpha > 0$ and q an integer such that $\alpha q \geq \eta_\varepsilon(R^\vee)$. Let $\mathbf{r} \in (\mathbb{R}^+)^n$ with $r_i \in [\alpha, \alpha']$ for all i . If \mathbf{s} is sampled from $D_{(1/q)R^\vee, \mathbf{r}}$, then $\|\mathbf{s}\|_{2, \infty} \leq \alpha' q \cdot \sqrt{d} \cdot \omega(\sqrt{\log N})$ with probability $\geq 1 - \varepsilon$.*

Proof. First, we know that $\|\mathbf{s}\|_{2, \infty} \leq \sqrt{d}\|\mathbf{s}\|_\infty$. Let $\varepsilon = N^{-\omega(1)}$, by assumption, we have that $\alpha q \geq \eta_\varepsilon(R^\vee)$. By Lemma 2.10 we know that $\|\mathbf{s}\|_\infty \leq \alpha' q \cdot \omega(\sqrt{\log N})$ with probability $\geq 1 - \varepsilon$. The result follows. □

In the following lemma, we transform a sample from $A_{q, \mathbf{s}, D_{(1/q)R^\vee, [\alpha, \alpha']}}^{(M)}$ to a sample of $A_{p, \mathbf{s}, \Psi_{\leq \beta}}^{(M)}$, assuming that $\|\mathbf{s}\|_{2, \infty}$ is bounded.

Lemma 4.26 (Adapted from [8, Le. 3.5]). *Let $\varepsilon = N^{-\omega(1)}$, $p, q > 2$, $\alpha, \alpha' \in (0, 1)$, and $s_{\max} > 0$. There is an efficient mapping from $R_q^d \times \mathbb{T}_{q, R^\vee}$ to $R_p^d \times \mathbb{T}_{p, R^\vee}$ which has the following properties:*

- *If the input is uniformly random, then the output is within negligible statistical distance from the uniform distribution.*
- *If the input is distributed from $A_{q, \mathbf{s}, D_{(1/q)R^\vee, [\alpha, \alpha']}}^{(M)}$, where $\mathbf{s} \in (R^\vee)^d$ with $\|\mathbf{s}\|_{2, \infty} \leq s_{\max}$, then the output distribution is within negligible statistical distance from $A_{p, \mathbf{s}', \Psi_{\leq \beta}}^{(M)}$, where \mathbf{s}' is uniform in $(R_p^\vee)^d$ and*

$$\beta^2 \geq 2 \left(\alpha' + \omega\left(\frac{p+q}{pq} s_{\max} \cdot \eta_\varepsilon(R^d)\right) \right).$$

Proof. The principle of this reduction is to first map $\mathbf{a} \in R_q^d$ to $\mathbf{a}' \in R_p^d$ by scaling it by p/q , and then randomly rounding it (using a discrete Gaussian distribution). Note that simply multiplying by p/q cannot work as, for example, the cardinality of R_p^d may not divide the cardinality of R_q^d . Then, we study the new error term, modified with the Gaussian rounding, and show that it is still a Gaussian error.

We sample \mathbf{s}_1 uniformly in $(R_p^\vee)^d$. On input $(\mathbf{a}, b) \in R_q^d \times \mathbb{T}_{q, R^\vee}$, the mapping is as follows:

- Set $\sigma = \omega((1 + p/q)\eta_\varepsilon(R^d))$,
- Sample \mathbf{d} from $D_{R^d - \frac{p}{q}\mathbf{a}, \sigma}$ and compute $\mathbf{a}' = \frac{p}{q}\mathbf{a} + \mathbf{d}$,
- Sample e_d from $D_{\sigma \cdot s_{\max}}$ and e' from $D_{\alpha'}$,
- Compute $b' = b + \frac{1}{p}\langle \mathbf{a}', \mathbf{s}_1 \rangle + \frac{1}{p}e_d + e'$,
- Return the new sample (\mathbf{a}', b') .

The choice of σ is derived from the proof of correctness of the reduction (see Lemmata 4.27 and 4.28 below).

We first show that the second step of the mapping transforms the uniform distribution in R_q^d to the uniform distribution in R_p^d , by considering the joint distribution of the pair $(\mathbf{a}', \mathbf{d})$. The following result can be interpreted as a simple particular case of [30, Se. 3].

Lemma 4.27. *Let $\varepsilon = N^{-\omega(1)}$ and assume that $\sigma \geq \omega((1 + \frac{p}{q})\eta_\varepsilon(R^d))$. Then the residual distribution of \mathbf{a}' is within negligible statistical distance to $U(R_p^d)$, and, for any $\bar{\mathbf{a}}' \in R_p^d$, the distribution of \mathbf{d} conditioned on $\mathbf{a}' = \bar{\mathbf{a}}'$ is within negligible statistical distance to $D_{\frac{p}{q}R^d + \bar{\mathbf{a}}', \sigma}$.*

Proof. Let $\bar{\mathbf{a}}' \in R_p^d$. Since $\mathbf{d} = \mathbf{a}' - \frac{p}{q}\mathbf{a} + p\mathbf{k}$ for some $\mathbf{k} \in R^d$ and $\mathbf{a} \in R_q^d$, we have that $\mathbf{d} - \mathbf{a}' \in \frac{p}{q}R^d$. Let $\bar{\mathbf{d}} \in \frac{p}{q}R^d + \bar{\mathbf{a}}'$. By construction, we have:

$$\Pr[\mathbf{a}' = \bar{\mathbf{a}}' \wedge \mathbf{d} = \bar{\mathbf{d}}] = \Pr\left[\mathbf{a} = \frac{q}{p}(\bar{\mathbf{a}}' - \bar{\mathbf{d}}) \wedge \mathbf{d} = \bar{\mathbf{d}}\right] = \frac{\rho_\sigma(\bar{\mathbf{d}})}{q^n \cdot \rho_\sigma(R^d - \bar{\mathbf{a}}' + \bar{\mathbf{d}})}.$$

In the latter, the denominator is within a factor $1 \pm \varepsilon$ from $q^n \cdot \rho_\sigma(R^d)$, because $\sigma \geq \eta_\varepsilon(R^d)$.

We now consider the residual distribution of \mathbf{a}' .

$$\begin{aligned} \Pr[\mathbf{a}' = \bar{\mathbf{a}}'] &= \sum_{\bar{\mathbf{d}} \in \frac{p}{q}R^d + \bar{\mathbf{a}}'} \Pr[\mathbf{a}' = \bar{\mathbf{a}}' \wedge \mathbf{d} = \bar{\mathbf{d}}] \\ &\in \frac{\rho_\sigma(\frac{p}{q}R^d + \bar{\mathbf{a}}')}{q^n \cdot \rho_\sigma(R^d)} \cdot [1 - \varepsilon, 1 + \varepsilon] \\ &\subseteq \frac{1}{p^n} \cdot [1 - \varepsilon, 1 + \varepsilon], \end{aligned}$$

because $\sigma \geq \eta_\varepsilon(\frac{p}{q}R^d)$.

Finally, we obtain that $\Pr[\mathbf{d} = \bar{\mathbf{d}} | \mathbf{a}' = \bar{\mathbf{a}}']$ is within a factor $1 \pm \varepsilon$ from a quantity that is proportional to $\rho_\sigma(\bar{\mathbf{d}})$. This completes the proof of the claim. \square

We now study the right hand size of the LWE sample. Assume that $b \in \mathbb{T}_{q, R^\vee}$ is of the form $\frac{1}{q}\langle \mathbf{a}, \mathbf{s} \rangle + e_q$ with $e_q \leftarrow D_{(1/q)R^\vee, [\alpha, \alpha']}$. Then we can write:

$$(\mathbf{a}', b') = \left(\mathbf{a}', \frac{1}{p}\langle \mathbf{a}', \mathbf{s} + \mathbf{s}' \rangle + \frac{1}{p}(\langle \mathbf{d}, \mathbf{s} \rangle + e_d) + e_q + e' \right). \quad (1)$$

The new error e_p is equal to $\frac{1}{p}(\langle \mathbf{d}, \mathbf{s} \rangle + e_d) + e_q + e'$. To study this new error, we first study the distribution of $\langle \mathbf{d}, \mathbf{s} \rangle + e_d$ conditioned on \mathbf{a}' in Lemma 4.28 (which generalizes [35, Co. 3.10] to the module case).

Lemma 4.28. *Let $s_{\max} > 0$ and $\mathbf{s} \in K^d$ with $\|\mathbf{s}\|_{2,\infty} < s_{\max}$. Let \mathbf{d} be distributed as $D_{(p/q)R^d-\mathbf{a},\sigma}$ for some arbitrary \mathbf{a} and $\sigma \geq \sqrt{2}(p/q)\eta_\varepsilon(R^d)$ and e be distributed as D_τ for some $\tau \geq \sigma \cdot s_{\max}$. Then the distribution of $\langle \mathbf{d}, \mathbf{s} \rangle + e$ is within negligible statistical distance of the elliptical Gaussian distribution $D_{\mathbf{t}}$ over K , where $t_i^2 = t_{\nu-i}^2 = \sigma^2 \sum_{k=1}^d |\sigma_i(s_k)|^2 + \tau^2$ for all i .*

Proof. By Lemma 2.13, we have that e is following the same distribution as $\langle \mathbf{e}_s, \mathbf{s} \rangle$ with \mathbf{e}_s distributed from $D_{\mathbf{r}', \dots, \mathbf{r}'}$ and $r'_i = r'_{\nu-i} = \tau / \sqrt{\sum_{k=1}^d |\sigma_i(s_k)|^2}$ for i .

As a consequence, we have that $\langle \mathbf{d}, \mathbf{s} \rangle + e$ is following the same distribution as $\langle \mathbf{d} + \mathbf{e}_s, \mathbf{s} \rangle$. We write $\mathbf{e}_s = \mathbf{e}_1 + \mathbf{e}_2$ with \mathbf{e}_1 distributed from $D_{\tau/S}$ and \mathbf{e}_2 distributed from $D_{(\sqrt{(r'_i)^2 - (\tau/s_{\max})^2})_i}$. We now use Lemma 2.8: As $\sigma \geq \sqrt{2}(p/q)\eta_\varepsilon(R^d)$ and $\tau \geq s_{\max} \cdot \sigma$, we have that $\mathbf{d} + \mathbf{e}_1$ is within statistical distance 4ε from $D_{\sqrt{\sigma^2 + (\tau/s_{\max})^2}}$. The quantity $\mathbf{d} + \mathbf{e}_s$ can be interpreted as the sum of two continuous Gaussians: It is within statistical distance 4ε from $D_{(\sqrt{\sigma^2 + (r'_i)^2})_i}$.

We use Lemma 2.13 once more. We obtain that $\langle \mathbf{d}, \mathbf{s} \rangle + e$ is within statistical distance 4ε from $D_{\mathbf{t}}$ with $t_i^2 = t_{\nu-i}^2 = \sigma^2 \sum_{k=1}^d |\sigma_i(s_k)|^2 + \tau^2$, for all i . \square

Let (\mathbf{a}, b) be sampled from $A_{q,\mathbf{s},D_{(1/q)R^\nu},[\alpha,\alpha']}^{(M)}$ and let (\mathbf{a}', b') be the image of (\mathbf{a}, b) by the mapping. To conclude the proof, we show that (\mathbf{a}', b') is sampled from $A_{p,\mathbf{s}',\Psi_{\leq\beta}}^{(M)}$:

- We recall that $b' = \frac{1}{p} \langle \mathbf{a}', \mathbf{s} + \mathbf{s}_1 \rangle + e_p$.
 - We showed that \mathbf{a}' is within negligible statistical distance from the uniform distribution in R_p^d .
 - We have that $\mathbf{s}' = \mathbf{s} + \mathbf{s}_1$, where \mathbf{s}_1 is uniform in $(R_p^\nu)^d$ and independent from \mathbf{s} . This ensures that $\mathbf{s}' \bmod p$ is uniform in $(R_p^\nu)^d$.
 - We now consider $e_p = \frac{1}{p} (\langle \mathbf{d}, \mathbf{s} \rangle + e_d) + e + e'$, where:
 - The component $\frac{1}{p} (\langle \mathbf{d}, \mathbf{s} \rangle + e_d)$ it is within negligible statistical distance from $D_{\mathbf{t}}$ with $t_i^2 = t_{\nu-i}^2 = \frac{1}{p} \sigma^2 \left(\sum_{k=1}^d |\sigma_i(s_k)|^2 + s_{\max}^2 \right)$ by applying Lemma 4.28.
 - The component $e + e'$ is within negligible statistical distance from $D_{\mathbf{t}'}$ with $(t'_i)^2 = (t'_{\nu-i})^2 = r_i^2 + (\alpha')^2$ by Lemma 2.8 and as, for all i , $\alpha' q \geq r_i q > \alpha q \geq \sqrt{2} \eta_\varepsilon(R^\nu)$.
- Then, the error component e_p is within negligible statistical distance from $D_{\mathbf{t}''}$ with $(t''_i)^2 = (t''_{\nu-i})^2 = r_i^2 + (\alpha')^2 + \frac{\sigma^2}{p^2} (\sum_{k=1}^d |\sigma_i(s_k)|^2 + s_{\max}^2)$. As $r_i \leq \alpha'$ holds for all i , and as $\|\mathbf{s}\|_{2,\infty} \leq s_{\max}$, we have:

$$t''_i = t''_{\nu-i} \leq \sqrt{2} \cdot \sqrt{(\alpha')^2 + \frac{\sigma^2}{p^2} s_{\max}^2} \leq \beta, \quad \text{for all } i.$$

\square

Reducing M-LWE $_{p,\Psi_{\leq\beta}}$ to M-LWE $_{p,\Upsilon_\beta}$. This is the last component of the proof of Theorem 4.8. The goal is to re-randomize the error distribution of M-LWE. The proof is adapted from [16, Le. 5.11].

Lemma 4.29. *Let $p \geq 2$ and $\beta \in (0, 1)$. There is a polynomial time reduction from M-LWE $_{p,\Psi_{\leq\beta}}$ to M-LWE $_{p,\Upsilon_\beta}$.*

Proof. Let $(\mathbf{a}, b = \frac{1}{p} \langle \mathbf{a}, \mathbf{s} \rangle + e)$ be a sample from $A_{p,\mathbf{s},D_{\mathbf{t}}}^{(M)}$ with $0 < t_i \leq \beta$ and $t_{\nu-i} = t_i$ for all i , and $\mathbf{s} \leftarrow U((R_q^\nu)^d)$. Let $(x'_i)_{i \in \mathbb{J}}$ be independent samples from $\Gamma(2, 1)$. We perform the following transformation:

$$(\mathbf{a}', b') := (\mathbf{a}, b + e'),$$

where e' is sampled from $D_{\mathbf{r}'}$, with \mathbf{r}' defined by $r'_i = r'_{\nu-i} = \beta^2 \sqrt{n} x'_i$ for all i .

This transformation maps the uniform distribution over $R_p^d \times \mathbb{T}_{R^\nu}$ to itself. On the other hand, it maps $A_{p,\mathbf{s},D_{\mathbf{t}}}^{(M)}$ to $A_{p,\mathbf{s},D_{\mathbf{r}'}}^{(M)}$, with $r'_i = r'_{\nu-i} = \sqrt{t_i^2 + \beta^2 \sqrt{n} x'_i}$, for all $i \in \mathbb{J}$.

Let S denote the set of ψ 's for which the oracle distinguishes with non-negligible probability between the uniform distribution over $R_p^d \times \mathbb{T}_{p,R^\nu}$ and the distribution $A_{p,\mathbf{s},\psi}^{(M)}$. By assumption, the measure of S under Υ_β is non-negligible. Lemma 2.2 implies that $D_{\mathbf{r}'} \in S$ with non-negligible probability. The result follows. \square

5 Converse reductions

In this section, we reduce both M-SIS and M-LWE to Mod-GIVP. This provides converse results to Theorems 3.6 and 4.7. We restrict the analysis to cyclotomic polynomials of the form $x^n + 1$ with n a power of 2, for the sake of simplicity. We expect the result to carry over to all cyclotomic polynomials, but this would add technical complications in the proof of Lemma 5.1. Choosing $x^n + 1$ implies that $R^\vee = \frac{1}{n}R$.

Reducing M-SIS to Mod-GIVP. Let $\mathbf{a}_1, \dots, \mathbf{a}_m$ be sampled uniformly and independently in R_q^d . Finding $\mathbf{z} = (z_1, \dots, z_m)^T \in R^m \setminus \{\mathbf{0}\}$ such that $\sum_i z_i \mathbf{a}_i = \mathbf{0} \pmod q$ and $\|\mathbf{z}\| \leq \beta$ corresponds to finding a short vector in the lattice:

$$\mathbf{A}^\perp = \{\mathbf{y} \in R^m : \mathbf{A}^T \mathbf{y} = \mathbf{0} \pmod q\},$$

where $\mathbf{A} \in R_q^{d \times m}$ is the matrix whose rows are the \mathbf{a}_i 's. As this lattice is a module lattice, if we solve Mod-GIVP $_{\gamma^\varepsilon}$ given as input an arbitrary basis of \mathbf{A}^\perp (which can be computed efficiently given \mathbf{A}), then we obtain a solution to the M-SIS instance, for $\beta = \gamma \cdot \eta_\varepsilon(\mathbf{A}^\perp)$. To assess the effectiveness of this reduction from M-SIS to Mod-GIVP, we are thus led to estimating $\eta_\varepsilon(\mathbf{A}^\perp)$ for \mathbf{A} sampled uniformly in $R_q^{m \times d}$. For this task, it is classical to study the dual lattice, as we have $\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2N(1+1/\varepsilon))}{\pi}} / \lambda_1^\infty(\Lambda^*)$ for any N -dimensional lattice Λ (see Lemma 2.5). The dual of the lattice \mathbf{A}^\perp is $\frac{1}{q}L_q(\mathbf{A})$ where

$$L_q(\mathbf{A}) = \{\mathbf{y} \in (R^\vee)^m : \exists \mathbf{s} \in (R^\vee)^d, \mathbf{B}\mathbf{s} = \mathbf{y} \pmod q\}.$$

Hence, it suffices to obtain a probabilistic lower bound on $\lambda_1^\infty(L_q(\mathbf{A}))$, for \mathbf{A} uniform in $R_q^{m \times d}$.

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

Lemma 5.1. *Let n, m, d, q be positive integers with $d \leq m$ and n a power of 2. We have:*

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

where $\lambda_1^{\infty, 2}(\cdot)$ refers to the lattice minimum with respect to $\|\cdot\|_{\infty, 2}$.

Proof. We generalize and adapt the proof of [37, Le. 8] (see also [38, Le. 3.2]). By the union bound, the probability that $L_q(\mathbf{A})$ contains a nonzero vector of infinity norm $\leq B := \frac{1}{8\sqrt{n}} q^{1 - \frac{d}{m}}$ is bounded from above by:

$$\sum_{\substack{\mathbf{t} \in (R_q^\vee)^m \\ 0 < \|\mathbf{t}\|_{\infty, 2} \leq B}} \sum_{\mathbf{s} \in (R_q^\vee)^d} \Pr_{\mathbf{A} \leftarrow U(R_q^{m \times d})} [\mathbf{A}\mathbf{s} = \mathbf{t}] = \sum_{\substack{\mathbf{t} \in (R_q^\vee)^m \\ 0 < \|\mathbf{t}\|_{\infty, 2} \leq B}} \sum_{\mathbf{s} \in (R_q^\vee)^d} \prod_{i \leq m} \Pr_{\mathbf{a} \leftarrow U(R_q^d)} [\mathbf{a}^T \mathbf{s} = t_i].$$

We now consider the probability (over the randomness of \mathbf{a}) that $\mathbf{a}^T \mathbf{s} = t_i$. For this purpose, we consider the decomposition of R_q as a Cartesian product of finite fields. By the Chinese Remainder Theorem, we know that $R_q \simeq R_q^\vee \simeq \mathbb{F}_{q^\delta} \times \dots \times \mathbb{F}_{q^\delta}$ for some integer δ dividing n (there are n/δ copies of the finite field of q^δ elements). Now, the equality $\mathbf{a}^T \mathbf{s} = t_i$ holds if and only if it holds over all n/δ CRT components. Wlog we consider the first one. If t_i and all the coordinates of \mathbf{s} are zero, then the probability is 1. Otherwise, if t_i or some coordinate of \mathbf{s} is nonzero on that first CRT component, then the probability is $\leq q^{-\delta}$. As a consequence, the probability under scope is bounded from above by:

$$\sum_{S \subseteq [n/\delta]} \sum_{\substack{\mathbf{s} \in (R_q^\vee)^d \\ \forall i, s_i \text{ is 0 on } S}} \sum_{\substack{\mathbf{t} \in (R_q^\vee)^m \\ 0 < \|\mathbf{t}\|_{\infty, 2} \leq B \\ \forall i, t_i \text{ is 0 on } S}} q^{m(|S|\delta - n)} \leq \sum_{S \subseteq [n/\delta]} \sum_{\substack{\mathbf{t} \in (R_q^\vee)^m \\ 0 < \|\mathbf{t}\|_{\infty, 2} \leq B \\ \forall i, t_i \text{ is 0 on } S}} q^{(m-d)(|S|\delta - n)}.$$

We now attempt to bound the number of t 's in R^\vee such that $0 < \|t\| \leq B$ and t is 0 on all CRT components corresponding to S . As $R^\vee = \frac{1}{n}R$, it suffices to bound the number of t 's in R such that $0 < \|t\| \leq nB$ and t is 0 on all CRT components corresponding to S .

The latter condition implies that t is a nonzero element of an ideal I of R of algebraic norm $q^{|S|\delta}$. Let $x \in I$ reaching $\lambda_1(I)$. By the arithmetic-geometric inequality, we have:

$$\lambda_1(I) = \|x\| \geq \sqrt{n}\mathcal{N}(x)^{1/n} = \sqrt{n}\mathcal{N}((x))^{1/n} \geq \sqrt{n}\mathcal{N}(I)^{1/n} = \sqrt{n}q^{|S|\delta/n}.$$

As a result, there is no such t when $|S| \geq (1 - d/m)n/\delta$. If $|S| \leq (1 - d/m)n/\delta$, then we are looking for the number of points of the (ideal) lattice I in the hyperball of radius nB and center 0. All such points are away from one another by at least $\lambda_1(I)$. Therefore, by the pigeon-hole principle, there are at most $(2nB/\lambda_1(I))^n \leq 4^{-n}q^{n - \frac{nd}{m} - |S|\delta}$ such points.

Now, the probability under scope can be bounded from above as

$$4^{-n} \cdot \sum_{S \subseteq [(1-d/m)n/\delta]} q^{(m-d)(|S|\delta-n)} \cdot q^{mn-nd-m|S|\delta} \leq 2^{-n}.$$

This completes the proof of the lemma. \square

As a consequence of the result above and the preceding discussion, we obtain the following converse to Theorem 3.6. 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 (the module rank in Mod-GIVP is m , which is possibly much larger than d).

Theorem 5.2. *For any $d \geq 1$ and $\varepsilon(N) = N^{-\omega(1)}$, there is a polynomial time reduction from solving M-SIS $_{q,m,\beta}$ to solving Mod-GIVP $_{q,\gamma^\varepsilon}$ (with module rank m), for any $m(N), q(N), \beta(N)$ and $\gamma(N)$ such that $\beta \geq \gamma\sqrt{N}\omega\left(\sqrt{\log(N/\varepsilon)}\right) \cdot q^{\frac{d}{m}}$ and $m, \log q \leq \text{poly}(N)$.*

Reducing M-LWE to Mod-GIVP. One of the classical ways for solving LWE consists in solving an associated SIS instance [25]. We propose an adaptation of this approach to module lattices: We reduce M-LWE to M-SIS and then combine this reduction with Theorem 5.2.

Let us sample \mathbf{s} uniformly in $(R_q^\vee)^d$, and ψ from \mathcal{Y}_α . More precisely, we sample x_i from $\Gamma(2, 1)$ for $i \in \mathbb{J}$, define $r_i = r_{\nu-i} = \alpha\sqrt{1 + \sqrt{n}x_i}$, and let $\psi = D_{\mathbf{r}}$. Assume that we have access to arbitrarily many samples $(\mathbf{a}_i, b_i) \in R_q^d \times \mathbb{T}_{R^\vee}$ with \mathbf{a}_i uniform in R_q^d and all the b_i 's uniform and independent in \mathbb{T}_{R^\vee} , or all the b_i 's of the form $b_i = \frac{1}{q}\langle \mathbf{a}_i, \mathbf{s} \rangle + e_i$ with the e_i 's are 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 $\mathbf{A} \in R_q^{m \times d}$ be the matrix whose rows are the \mathbf{a}_i 's. By solving M-SIS $_{q,m,\beta}$ for \mathbf{A}^T , we obtain a nonzero vector $\mathbf{z} \in R^m$ such that $\|\mathbf{z}\| \leq \beta$ and $\mathbf{z}^t \cdot \mathbf{A} = \mathbf{0} \pmod{q}$. Now, we compute $\langle \mathbf{z}, \mathbf{b} \rangle$, where $\mathbf{b} \in \mathbb{T}_{R^\vee}^m$ is the vector made of the b_i 's. If the b_i 's are uniform independent of the \mathbf{a}_i 's, then the inner product $\langle \mathbf{z}, \mathbf{b} \rangle$ is uniformly distributed in \mathbb{T}_{R^\vee} . Otherwise, we have $\langle \mathbf{z}, \mathbf{b} \rangle = \langle \mathbf{z}, \mathbf{e} \rangle \pmod{R^\vee}$, where \mathbf{e} is the vector made of the e_i 's. By Lemma 2.13, we have that $\langle \mathbf{z}, \mathbf{e} \rangle$ is distributed as $D_{\mathbf{r}'}$ with $r'_j = r_j \cdot \sqrt{\sum_{k \leq m} |\sigma_j(z_k)|^2}$ for all $j \in \mathbb{Z}_\nu^\times$. As a consequence, we have

$$\begin{aligned} \|\langle \mathbf{z}, \mathbf{b} \rangle\| &\leq t\sqrt{n} \cdot \max_j |r'_j| \\ &\leq t\sqrt{n} \cdot \|\mathbf{z}\| \cdot \max_j |r_j| \leq 2tn^{3/4}\alpha\beta \cdot \max_j |x_j|, \end{aligned}$$

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

Overall, we have proved that if β is such that $n^{7/4}\omega(\log(N)) \cdot \alpha\beta < 1$, then we can distinguish between the two challenge distributions with non-negligible advantage. By Theorem 5.2, we thus obtain a reduction from $\text{Mod-GIVP}_\gamma^{\eta_\varepsilon}$ with module rank m to $\text{M-LWE}_{q,r_\alpha}$, if γ is such that $\alpha\gamma n^{7/4}\sqrt{N}\omega(\sqrt{\log(N/\varepsilon)})q^{d/m} < 1$. Taking $m = d \log q$ leads to the following result.

Theorem 5.3. *For any $d \geq 1$ and $\varepsilon(N) = N^{-\omega(1)}$, there is a probabilistic polynomial time reduction from solving $\text{M-LWE}_{q,r_\alpha}$ to solving $\text{Mod-GIVP}_\gamma^{\eta_\varepsilon}$ (with module rank $d \log q$), for any $\alpha(N)$ and $\gamma(N)$ such that $\frac{1}{\alpha} \geq \gamma N^{3/2}\omega(\sqrt{\log(N/\varepsilon)})$ and $\log q \leq \text{poly}(N)$.*

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 and Monash Universities, whose hospitalities are 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. J. Alperin-Sheriff and C. Peikert. Circular and KDM security for identity-based encryption. In *Proc. of PKC*, volume 7293 of *LNCS*, pages 334–352. Springer, 2012.
3. 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.
4. 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.
5. 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.
6. W. Bosma and M. Pohst. Computations with finitely generated modules over Dedekind rings. In *Proc. of ISSAC*, pages 151–156, 1991.
7. Z. Brakerski, C. Gentry, and V. Vaikuntanathan. Fully homomorphic encryption without bootstrapping. *Cryptography ePrint Archive, Report 2011/277*, 2011.
8. Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *Proc. of STOC*, pages 575–584. ACM, 2013.
9. H. Cohen. *Advanced topics in computational number theory*. Springer, 2000.
10. 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.
11. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008.
12. 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.
13. J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In *Proc. of ANTS-III*, pages 267–288, 1998.
14. 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.
15. 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.
16. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *Proc. of EUROCRYPT*, LNCS, pages 1–23. Springer, 2010. All result numberings used in the present article correspond to those of the draft of the full version, available at <http://eprint.iacr.org/2012/230.pdf>.
17. V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for Ring-LWE cryptography. In *Proc. of EUROCRYPT*, volume 7881 of *LNCS*, pages 35–54. Springer, 2013.
18. 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 [20].

19. 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.
20. D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complexity*, 16(4):365–411, 2007. Full version of [18].
21. D. Micciancio and S. Goldwasser. *Complexity of lattice problems: a cryptographic perspective*. Kluwer Academic Press, 2002.
22. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller, 2011. IACR Cryptology ePrint Archive, report 2011/501.
23. 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 [24].
24. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Full version of [23].
25. 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.
26. R. A. Mollin. *Algebraic Number Theory*. Chapman and Hall/CRC Press, 1999.
27. A. O'Neill, C. Peikert, and B. Waters. Bi-deniable public-key encryption. In *Proc. of CRYPTO*, volume 6841 of *LNCS*, pages 525–542. Springer, 2011.
28. C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Comput. Complexity*, 2(17):300–351, 2008.
29. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.
30. C. Peikert. An efficient and parallel gaussian sampler for lattices. In *Proc. of CRYPTO*, volume 6223 of *LNCS*, pages 80–97. Springer, 2010.
31. 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.
32. 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.
33. K. Pietrzak. Subspace LWE, 2011. Available at <http://homepages.cwi.nl/~pietrzak/publications/SLWE.pdf>.
34. O. Regev. Lecture notes of *lattices in computer science*, taught at the Computer Science Tel Aviv University. Available at http://www.cims.nyu.edu/~regev/teaching/lattices_fall_2009/index.html.
35. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proc. of STOC*, pages 84–93, 2005.
36. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009. Full version of [35].
37. 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.
38. D. Stehlé and R. Steinfeld. Making NTRUEncrypt and NTRUSign as secure as standard worst-case problems over ideal lattices. Cryptology ePrint Archive, Report 2013/004, 2013. Full version of [37].
39. 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.