Hardness of Decision (R)LWE for Any Modulus

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

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

Abstract. The decision Learning With Errors problem has proven an extremely flexible foundation for devising provably secure cryptographic primitives. LWE can be expressed in terms of linear algebra over $\mathbb{Z}/q\mathbb{Z}$. This modulus q is the subject of study of the present work. When q is prime and small, or when it is exponential and composite with small factors, LWE is known to be at least as hard as standard worst-case problems over euclidean lattices (sometimes using quantum reductions). The Ring Learning With Errors problem is a structured variant of LWE allowing for more compact keys and more efficient primitives. It is known to be at least as hard as standard worst-case problems restricted to so-called ideal lattices, but under even more restrictive arithmetic conditions on q.

In this work, we prove that the arithmetic form of the modulus q is irrelevant to the computational hardness of LWE and RLWE. More precisely, we show that these problems are at least as hard as standard worst-case problems on lattices, under the unique condition that q is of polynomial bit-size. This result is most useful for adapting LWE-based cryptographic constructions to the RLWE setting. Among others, this allows us to derive the first Identity-Based Encryption scheme of quasi-optimal performance proven secure under standard worst-case lattice assumptions, in the standard model. Other applications include authentication, functional encryption and traitor tracing.

Keywords. LWE, RLWE, lattice-based cryptography, provable security.

1 Introduction

The decision Learning With Errors problem, introduced by Regev in [37,38] (see also the survey [39]), has proven an invaluable tool for designing provably secure cryptographic protocols. These cryptographic functionalities include, among others, IND-CPA public-key encryption [37,38,19], IND-CCA public-key encryption [33], (hierarchical) identity-based encryption [19,14,1,2], functional encryption [4,3], traitor tracing [25], group signature [20], identification [23], lossy trapdoor function [35] and homomorphic encryption [18,12]. LWE consists in distinguishing from uniform arbitrarily many samples of the form $(\boldsymbol{a}_i, \langle \boldsymbol{a}_i, \boldsymbol{s} \rangle + e_i) \in \mathbb{Z}_q^{n+1}$, where the vectors $\boldsymbol{a}_i$ are sampled independently and uniformly in $\mathbb{Z}_q^n$, the vector $\boldsymbol{s}$ is sampled uniformly but common to all pairs, and the errors e_i are sampled independently from a variant of the Gaussian distribution with standard deviation αq that is small relative to q (as otherwise the problem becomes vacuously hard). The parameters α and q are functions of n, which we let grow to infinity when studying the hardness of LWE. An extremely attractive feature of LWE is that under some conditions on α and q, it is proven at least as hard as standard worst-case problems over euclidean lattices. For instance, Regev [37,38] gave a (quantum) reduction from the approximate Shortest Independent Vectors Problem (SIVP) to LWE, when q is polynomial in n and prime (and some other conditions). Peikert [33] gave a classical reduction for q exponential in n and product of many distinct prime factors. The conditions on q for these reductions have been relaxed further in [7,30], to handle composite moduli involving only small prime factors.

The decision Ring Learning With Errors problem, introduced by Lyubashevsky et al in [27], allows for more compact and more efficient cryptographic constructions than LWE. It consists in distinguishing from uniform arbitrarily many samples of the form $(a_i, a_i \cdot s + e_i) \in R_q^2$, where R_q is typically the ring $\mathbb{Z}_q[x]/(x^n + 1)$ with n a power of 2, the ring elements a_i are sampled independently and uniformly, the element s is sampled uniformly but common to all pairs, and the e_i's are sampled independently from a variant of the n-dimensional Gaussian distribution and interpreted as ring elements, with a standard deviation αq that is small relative to q. The efficiency improvement stems from the fact that an RLWE sample encodes n (non-independent) LWE samples at once. Cryptographic applications of RLWE include fast encryption [27,41], fast lossy trapdoor functions [43], and fast homomorphic encryption [13,11,17]. RLWE is known [27] to be at least as hard as approximate SIVP restricted to the class of so-called ideal lattices, when q is polynomial in n, prime, and such that x^n+1 has n distinct linear factors modulo q. These conditions on the modulus q are significantly more restrictive than their LWE counterparts.

The strong restrictiveness of the modulus conditions of RLWE relative to those of LWE prevents one from extending in a straightforward manner a number of LWE-based cryptographic constructions to the RLWE setting. For example, in the (hierarchical) identity-based encryption ((H)IBE) scheme from [1], the identities are encoded into full rank difference matrices (i.e., any difference of two distinct matrices in the set is required full-rank). These are typically implemented by encoding in matrix format a polynomial ring $\mathbb{Z}_q[x]/\Phi(x)$ such that Φ is irreducible modulo q. Note that in this case, the ring $\mathbb{Z}_q[x]/\Phi(x)$ is isomorphic to the finite field $\mathbb{F}_{q^n}$. This algebraic structure seems incompatible with that imposed by the RLWE restrictive modulus conditions, which force the ring R_q to be isomorphic to the ring $(\mathbb{F}_q)^n$ containing plentiful divisors of zero (extending the proof from [1] fails because the simulator does not seem able to handle these non-divisors of zero). In a number of other LWE-based constructions (such as [36,23,25]), LWE is interpreted as a problem of (noisy) linear algebra over $\mathbb{Z}_q$. When q is prime, the ring $\mathbb{Z}_q$ is a field, and one can exploit properties that hold on vector spaces. The restrictive RLWE conditions force R_q to be very far from a field, and the vector spaces degenerate to modules over the non-field ring R_q. For instance, the property that is central to the proof of [36] and fails to hold for such a module is that if the columns of a matrix M are linearly independent, then the map $\boldsymbol{x} \mapsto M \cdot \boldsymbol{x}$ is onto.

All known hardness proofs for decision (R)LWE [38,33,29,27,30] proceed by reducing the computational variant of (R)LWE, which is proven hard under much milder modulus conditions, to the decisional variant. These computational to decisional reductions rely on guessing the secret $\boldsymbol{s}$ common to all samples. As the range of $\boldsymbol{s}$ is of size q^n, the guess is broken into smaller guesses, based on the decomposition of the range of $\boldsymbol{s}$ into a Cartesian product of small finite fields. For example, Regev [38] set q prime and polynomial in n so that each component of $\mathbb{Z}_q^n$ is a small field. Peikert [33] handled a larger q by taking it very composite, in order to split $\mathbb{Z}_q$ into smaller pieces via the Chinese Remainder Theorem. Lyubashevsky et al [27] chose q small such that x^n+1 has n distinct linear factors modulo q, in order to break R_q into a Cartesian product of n copies of $\mathbb{F}_q$. These approaches seem bound to fail for moduli q such that $\mathbb{Z}_q$ and R_q contain large subrings that are fields.

Our main result. We show that (decision) LWE and RLWE are at least as hard as standard worst-case problems over lattices (restricted to ideal lattices in the case of RLWE), without any condition on the arithmetic form of q. The only remaining constraint on q is that it can be written in polynomial time (i.e., $q \leq 2^{n^{O(1)}}$). For this purpose, we introduce a modulus-switching self-reduction for (R)LWE: For specific choices of the noise distributions, $\mathrm{LWE}_{q,\alpha}$ (resp. $\mathrm{RLWE}_{q,\alpha}$) can be reduced to $\mathrm{LWE}_{p,\beta}$ (resp. $\mathrm{RLWE}_{p,\beta}$) if:[4]

$$\beta \geq \alpha \cdot \max\left(1, \frac{q}{p}\right) \cdot \widetilde{\Omega}\left(n^{1/2}\right) \quad \text{and} \quad \alpha q = \widetilde{\Omega}(1),$$

$$\text{resp.} \quad \beta \geq \alpha \cdot \max\left(1, \frac{q}{p}\right) \cdot \widetilde{\Omega}\left(n^{3/4}\right) \quad \text{and} \quad \alpha q = \widetilde{\Omega}(n^{1/2}).$$

The conditions on αq are typically much weaker than requirements on αq for the worst-case to average-case reductions to hold. This self-reduction can be combined with the known hardness proofs for specific choices of q [38,33,27] to obtain classical/quantum reductions from standard (ideal) lattice problems. This provides the first hardness results for LWE with a modulus q that is prime and large, for RLWE with a modulus q that is a power of 2 (which could be an interesting choice from an implementation perspective), and for RLWE with a modulus q such that x^n+1 has few factors modulo q.

Cryptographic applications. Our cryptographic applications all derive from the hardness of RLWE with q prime such that x^n+1 (with n a power of 2) has exactly two irreducible factors modulo q. Note

[4] The $\widetilde{\Omega}$ notation absorbs some factor that is polynomial in $\log n$.

that ideally, we would like only one factor, but such a modulus q does not exist. The closest that can be achieved is to take q prime with $q = 3 \bmod 8$, since in that case the polynomial $x^n + 1$ has exactly two distinct irreducible factors of degrees $n/2$. This provides a ring isomorphism $R_q \simeq (\mathbb{F}_{q^{n/2}})^2$, thus illustrating that zero divisors of such an R_q are extremely uncommon. For this reason, the ring R_q "behaves almost like a finite field," and modules over R_q "behave almost like vector spaces over a finite field."

As a first application, we can encode the identities of the (H)IBE scheme from [1] into the diagonal of R_q (the set of elements of R_q with equal $\mathbb{F}_{q^{n/2}}$-components), leading to a family of full-rank difference matrices that both supports an exponential number of identities and has an algebraic structure compatible with R_q. In algebraic terms, the security proof requires an additive subgroup of matrices over R_q where any non-zero element is invertible (i.e., a division algebra over R_q), whereas the efficiency of the resulting scheme is polynomial in the dimension of those matrices: The CRT diagonal of R_q provides such a one-dimensional division algebra. This adaptation of [1] to the RLWE setting leads to the first asymptotically quasi-optimal IBE (key sizes quasi-linear in the security parameter λ, and encryption/decryption of λ bits costing $\widetilde{O}(\lambda)$ bit operations) that is provably secure under a standard lattice assumption, in the standard model. This efficiency improvement also applies to the revocable IBE scheme from [15] and to the completely non-malleable encryption scheme from [40], which both build upon [1]. By relying on the construction from [9], this also provides an asymptotically quasi-optimal IND-CCA encryption scheme (proven secure in the standard model). Note that this was already possible via lossy trapdoor functions [35,43]. Finally, our work also leads to a RLWE-based variant of the functional encryption scheme for inner product predicates from [4], which also builds upon [1]. This is most useful for its applications, such as hidden vector encryption [10] or predicate encryption for CNF/DNF formulæ [22], that require a scalar domain of exponential size (so that a uniformly chosen scalar is zero with exponentially small probability). In this context, moving from LWE to RLWE allows one to decrease the sizes of the public parameters ciphertexts from $\widetilde{O}(\ell\lambda^5)$ and $\widetilde{O}(\ell\lambda^4)$ to $\widetilde{O}(\ell\lambda)$ and $\widetilde{O}(\ell\lambda)$ respectively, where ℓ is the length of the predicate inner products. Finally, the largeness of the CRT diagonal of R_q also leads to a RLWE variant of the authentication scheme from [21].

Another generic application of our result is to extend to the RLWE setting the LWE-based constructions that were exploiting results holding for vector spaces. By choosing q such that R_q is isomorphic to $(\mathbb{F}_{q^{n/2}})^2$, the modules over R_q behave almost like vector spaces. For example, if the columns of the matrix M over R_q are uniform conditioned on being linearly independent, then the map $\boldsymbol{x} \mapsto M \cdot \boldsymbol{x}$ is onto with probability exponentially close to 1. This provides a RLWE variant of the Subspace LWE problem from [36]. This also leads to an asymptotically efficient RLWE variant of the traitor tracing scheme from [25].

Overview of the technique. The modulus-switching self-reduction of (R)LWE proceeds in several steps. To fix the ideas, assume we aim at reducing $\mathrm{LWE}_{q,\alpha}$ to $\mathrm{LWE}_{p,\beta}$.

- *Discretizing the noise.* We start from a variant of $\mathrm{LWE}_{q,\alpha}$ where the noise distribution is a discrete Gaussian with support $\mathbb{Z}$. From [20], this variant is no easier than the original LWE problem from [38]. Discretizing the noise distribution allows us to use the HNF variant of LWE (see just below), and using a discrete Gaussian rather than a rounded continuous Gaussian facilitates the noise handling of the $\mathrm{LWE}_{p,\beta}$ samples (see 'handling the RHS' below).
- *Replacing* LWE *by* HNF-LWE. Our next observation is that we can use $\mathrm{LWE}_{q,\alpha}$ samples where the coordinates of the secret $\boldsymbol{s} \in \mathbb{Z}_q^n$ is chosen from the noise distribution instead of the uniform distribution on $\mathbb{Z}_q^n$. This HNF-ization, introduced in [7], is crucial here for limiting the amplification of the noise in the Right Hand Size (RHS) of the LWE samples, in the modulus-switching step.
- *Switching the modulus.* So far, we have either samples from the uniform distribution over $\mathbb{Z}_q^{n+1}$ or $\mathrm{LWE}_{q,\alpha}$ samples $(\boldsymbol{a}_i, \langle \boldsymbol{a}_i, \boldsymbol{s}\rangle + e_i)$ with $\boldsymbol{s}$ small and the e_i's sampled from a discrete Gaussian distribution of standard deviation αq. We aim at mapping these samples to either uniform samples over $\mathbb{Z}_p^{n+1}$ or $\mathrm{LWE}_{p,\beta}$ samples. We proceed by refining the modulus-switching technique from [12,13], introduced in the context of fully homomorphic encryption. We multiply the components by p/q and "round" them to integers. To ensure the resulting Left Hand Sides (LHS) are uniformly distributed over $\mathbb{Z}_p^n$, we rely on Peikert's convolution technique [34]. More precisely, we replace $\boldsymbol{a}_i$ by a discrete Gaussian sample $\boldsymbol{r}_i$ with support $\mathbb{Z}^n$, small standard deviation and center $\frac{p}{q}\boldsymbol{a}_i$.

- *Handling the RHS.* The randomized rounding of the previous step induces an additional error term $\langle \boldsymbol{r}_i, \boldsymbol{s} \rangle$ in the RHS. The smallness of both $\boldsymbol{r}_i$ and $\boldsymbol{s}$ ensures that this new error term is small. By combining the former error term, this new error term and a continuous Gaussian error term, we obtain a resulting noise distribution for the $\mathrm{LWE}_{q,\alpha}$ samples that is a continuous Gaussian with standard deviation chosen arbitrarily in a specific interval related to β.
- *Re-randomizing the noise distribution.* To allow for a fully average-case variant of $\mathrm{LWE}_{p,\beta}$, we finally randomize the noise distribution, using the same technique as in [27, Se. 5].

Related works. Thanks to its cryptographic attractiveness, the hardness of (R)LWE has been investigated in a number of works. Different noise distributions have been shown to allow for hardness based on standard lattice assumptions [38,33,20,27,8]. (R)LWE was shown to remain hard even for a secret $\boldsymbol{s}$ that is derived from the noise distribution [7]. The sample complexity of the hardness proof of decision LWE was considered in [29] (note that all the steps of our modulus-switching self-reduction preserve the number of LWE samples, except the HNF-ization step). An extension of (R)LWE to a module setting generalizing both LWE and RLWE was proposed in [11] and proven to admit a reduction from lattice problems in [24]. The hardness of (R)LWE with respect to the modulus q was considered in [38,33,7,30,27]. Our technique allows one to handle much more moduli, at the expense of increasing the approximation factors of the worst-case lattice problems by factors $\widetilde{\Omega}(n^{1/2})$ and $\widetilde{\Omega}(n^{3/4})$ for LWE and RLWE respectively.

Open problems. An original motivation for this work was to fully dequantumize Regev's reduction from standard lattice problems to LWE [38]. This was achieved by Peikert [33], but only for moduli that are exponentially large and products of many small primes. By using Peikert's result for such a special modulus q, and then switching to another modulus p, one obtains a classical reduction from standard lattice problems to LWE with that new modulus p. This would be most interesting for a modulus p that is as small as those handled by Regev's quantum reduction. However, applying our reduction with these parameters leads to a standard deviation β that is exponentially larger than α. As the problem becomes vacuous for $\beta = \widetilde{\Omega}(1)$, this means that α should be chosen exponentially small, but then $\mathrm{LWE}_{q,\alpha}$ can be solved in polynomial-time.

Another natural open problem would be to remove the increase of the standard deviation by factors $\widetilde{\Omega}(n^{1/2})$ and $\widetilde{\Omega}(n^{3/4})$, when switching moduli. This increase limits the range of moduli p for which we obtain non-vacuous hardness results based on standard lattice problems. To fix the ideas, consider $\mathrm{LWE}_{p,\beta}$. As it is vacuously hard for $\beta = \widetilde{\Omega}(1)$ and Regev's hardness proof for $\mathrm{LWE}_{q,\alpha}$ assumes that $\alpha q = \widetilde{\Omega}(n^{1/2})$, the standard deviation increase implies that the LWE self-reduction provides non-vacuous results only for $p \geq n$.

Road-map. In Section 2, we introduce the necessary background for the exposition of our results. In particular, we recall the module variant of LWE, to place ourselves in an algebraic setup that allows us to handle both LWE and RLWE at once (and thus avoid cumbersome duplications of the proofs). Section 3 contains the modulus-switching self-reduction. Finally, we discuss cryptographic applications in Section 4.

2 Preliminaries

Notations. We use standard Landau notations. A function $f(n)$ is said negligible (resp. exponentially small) if $f(n) = n^{-\omega(1)}$ (resp. $f(n) = 2^{-\Omega(n)}$). The statistical distance between two distributions X and Y on a countable set D is defined as follows: $\Delta(X, Y) = \frac{1}{2}\sum_{d \in D} |X(d) - Y(d)|$. For a set A with finite measure, we let $U(A)$ denote the uniform distribution on A. All vectors will be denoted in bold. The hermitian norm of a vector $\boldsymbol{x} \in \mathbb{C}^n$ will be denoted by $\|\boldsymbol{x}\|$. Similarly, its infinity norm $\max_i |x_i|$ will be denoted by $\|\boldsymbol{x}\|_\infty$. If two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ over the same ring share the same number of coordinates, their inner product will be denoted by $\langle \boldsymbol{x}, \boldsymbol{y} \rangle$. For any vectors $\boldsymbol{c}, \boldsymbol{r} \in \mathbb{R}^n$ with $r_i > 0$ for all i, we define the function $\rho_{\boldsymbol{r},\boldsymbol{c}}(\boldsymbol{x}) = \exp(-\pi \sum_{i \leq n} \frac{(x_i - c_i)^2}{r_i^2})$, for $\boldsymbol{x} \in \mathbb{R}^n$. We extend this to any countable set $A \subseteq \mathbb{R}^n$ in the usual way: $\rho_{\boldsymbol{r},\boldsymbol{c}}(A) = \sum_{\boldsymbol{x} \in A} \rho_{\boldsymbol{r},\boldsymbol{c}}(\boldsymbol{x})$. If the vector $\boldsymbol{r}$ is constant equal to $\sigma > 0$, we write $\rho_{\sigma,\boldsymbol{c}}$ instead.

2.1 Algebraic setup

A cyclotomic number field. Let $n \geq 1$ be a power of 2 and $K = \mathbb{Q}(\xi)$ with $\xi = \exp(i\pi/n)$ denote the $(2n)$-th cyclotomic number field.[5] The field K is a degree n extension of $\mathbb{Q}$, and we let $(\sigma_j)_{j \leq n}$ denote the canonical embeddings, ordered so that $\sigma_{j+n/2} = \overline{\sigma_j}$ for all $j \leq n/2$. For $y \in K$, the notation $\sigma_C(y)$ will refer to the vector $(\sigma_j(y))_j \in \mathbb{C}^n$. As in [27], we will consider the subspace $H = \{(x_1, \ldots, x_n) \in \mathbb{C}^n : \forall j \leq n/2, x_{n/2+j} = \overline{x_j}\}$, which is isomorphic to $K_\mathbb{R} = K \otimes_\mathbb{Q} \mathbb{R}$. Let $\boldsymbol{h}_j = \frac{1}{\sqrt{2}}(\boldsymbol{e}_j + \boldsymbol{e}_{j+n/2})$ and $\boldsymbol{h}_{j+n/2} = \frac{i}{\sqrt{2}}(\boldsymbol{e}_j - \boldsymbol{e}_{j+n/2})$ for $j \leq n/2$, where $\boldsymbol{e}_j$ denotes the j-th canonical basis vector of $\mathbb{C}^n$. Then the $\boldsymbol{h}_j$'s form a basis of H as a real vector space. For any $y \in K$, there exists a unique $(y_j)_j \in \mathbb{R}^n$, which we denote by $\sigma_H(y)$, such that $\sigma_C(y) = \sum_j y_j \boldsymbol{h}_j$.

A ring and its ideals. We let $R = \mathbb{Z}[\xi]$ denote the ring of integers of K. It is isomorphic to the polynomial ring $\mathbb{Z}[x]/(x^n + 1)$, and an explicit isomorphism can be derived by deleting rows and columns from the $(2n)$-dimensional complex Fourier transform (we refer to [41] for more details). We let $R^* = \{x \in K : \forall y \in R, \langle \sigma_C(x), \sigma_C(y) \rangle \in \mathbb{Z}\}$ denote the dual of R. For our particular choice of K, we have $R^* = \frac{1}{n}R$. Any additive subgroup I of R that is closed under multiplication by every element of R is called an ideal of R. The set of the equivalence classes $g + I$ of R modulo I is denoted by R/I. If x is a ring element, we let (x) denote the ideal generated by x. For $q > 2$ prime, the ring $R_q = R/(q)$ is isomorphic to the Cartesian product of finite fields $(\mathbb{F}_{q^k})^{\frac{n}{k}}$ where $k < n$ is the power of 2 that is the common degree of all irreducible factors of $x^n + 1$ modulo q. For $q = 1 \bmod 2n$, we have $k = 1$, and for $q = 3 \bmod 8$, we have $k = n/2$.

Modules. Let $d \geq 1$. A subset $M \subseteq K^d$ is called an R-module if it closed under addition and multiplication by elements of R. It is a finitely generated module if there exists a finite family $(\boldsymbol{b}_i)_i$ such that $M = \sum_i R \cdot \boldsymbol{b}_i$. We say that a module $M \subseteq K^d$ is full-rank if its K-span has K-dimension d. For the sake of simplicity, we will restrict ourselves to full-rank modules. We extend σ_H to modules, by considering the concatenation of the embeddings of the successive coordinates. We define the euclidean (resp. infinity) norm of a vector $\boldsymbol{x} \in K^d$ by $\|\boldsymbol{x}\| = \|\sigma_H(\boldsymbol{x})\|$ (resp. $\|\boldsymbol{x}\|_\infty = \|\sigma_H(\boldsymbol{x})\|_\infty$). We will also use the mixed norm $\|\boldsymbol{x}\|_{2,\infty} = \max_{j \leq n} \sqrt{\sum_{k \leq d} |\sigma_j(x_k)|^2}$. We have $\|\boldsymbol{x}\|_{2,\infty} \leq \sqrt{d}\|\boldsymbol{x}\|_\infty$ for all $\boldsymbol{x} \in K^d$.

Lattices. A euclidean lattice $\mathit{\Lambda}$ is a discrete additive subgroup of $\mathbb{R}^n$. It can be written $\mathit{\Lambda} = \sum_i \mathbb{Z}\boldsymbol{b}_i$ for some linearly independent vectors $(\boldsymbol{b}_i)_{i \leq k} \in \mathbb{R}^n$, which we call a basis of $\mathit{\Lambda}$. For the sake of simplicity, we will restrict ourselves to full rank lattices, i.e., with $k = n$. For an n-dimensional lattice $\mathit{\Lambda}$ and for any $i \leq n$, we define the i-th minimum $\lambda_i(\mathit{\Lambda})$ as the smallest r such that $\mathit{\Lambda}$ contains $\geq i$ linearly independent vectors of (euclidean) norm $\leq r$. The dual lattice of $\mathit{\Lambda} \subseteq \mathbb{R}^n$ is defined as $\mathit{\Lambda}^* = \{\boldsymbol{x} \in \mathbb{R}^n : \forall \boldsymbol{y} \in \mathit{\Lambda}, \langle \boldsymbol{x}, \boldsymbol{y} \rangle \in \mathbb{Z}\}$. If I is a non-zero ideal of R, then $\sigma_H(I)$ is an n-dimensional lattice. We call such lattices ideal lattices. Similarly, if $M \subseteq K^d$ is a full-rank module, then $\sigma_H(M)$ is an nd-dimensional lattice. We call such lattices (n,d)-module lattices. Note that $(n,1)$-module lattices contained in $\mathbb{Z}^n$ are exactly ideal lattices, and that $(1,n)$-module lattices are arbitrary n-dimensional lattices.

2.2 Gaussian measures

For $\sigma > 0$, we let ν_σ denote the n-dimensional continuous Gaussian distribution of standard deviation σ: we have $\nu_\sigma(\boldsymbol{x}) = \rho_\sigma(\boldsymbol{x})/\sigma^n$ for any $\boldsymbol{x} \in \mathbb{R}^n$. We extend it to elliptical Gaussian distributions: For $\boldsymbol{r} = (r_i)_{i \leq n}$ with $r_i > 0$ for all i, the distribution $\nu_{\boldsymbol{r}}$ is defined as $\nu_{\boldsymbol{r}} = (\nu_{r_i})_{i \leq n}$. We define $\mathit{\Psi}_{[\alpha,\alpha']}$ for $0 \leq \alpha < \alpha'$, as the set of Gaussian distributions $\nu_{\boldsymbol{r}}$ with $\alpha < r_i \leq \alpha'$, for all i. We write $\mathit{\Psi}_{\leq \alpha'}$ when $\alpha = 0$. We will also use the distribution $\mathit{\Upsilon}_\alpha$ introduced in [28]. The gamma distribution $\mathit{\Gamma}(2,1)$ has density $x \exp(-x)$ for $x \geq 0$ and zero for $x < 0$. For $\alpha > 0$, a distribution sampled from $\mathit{\Upsilon}_\alpha$ is an elliptical Gaussian distribution $\nu_{\boldsymbol{r}}$ whose parameters are $r_i = r_{i+n/2} = \alpha\sqrt{(1 + \sqrt{n}x_i)}$ for $i \leq n/2$, where the x_i's are sampled independently from the distribution $\mathit{\Gamma}(2,1)$. We will use the following result from [28].

Lemma 1 ([28, Claim 5.10]). *Let P be the distribution $\mathit{\Gamma}(2,1)^n$ and Q be the distribution $(\mathit{\Gamma}(2,1) - z_1) \times \ldots \times (\mathit{\Gamma}(2,1) - z_n)$ for some $0 \leq z_1, \ldots, z_n \leq 1/\sqrt{n}$. Then for any measurable set $A \subseteq \mathbb{R}^n$, we have $\int_A Q \geq \frac{1}{poly(n)} \cdot (\int_A P)^2$.*

[5] Our techniques should apply to all cyclotomic fields, but we restrict ourselves to this setting, for simplicity.

For any $\boldsymbol{c}, \boldsymbol{r} \in \mathbb{R}^n$ with $r_i > 0$ for all i and for any n-dimensional lattice Λ, the discrete Gaussian distribution of support Λ, standard deviation $\boldsymbol{r}$ and center $\boldsymbol{c}$ is defined by $D_{\Lambda,\boldsymbol{r},\boldsymbol{c}}(\boldsymbol{x}) = \frac{\rho_{\boldsymbol{r},\boldsymbol{c}}(\boldsymbol{x})}{\rho_{\boldsymbol{r},\boldsymbol{c}}(\Lambda)}$. If $\boldsymbol{r}$ is constant equal to $\sigma > 0$, we write $D_{\Lambda,\sigma,\boldsymbol{c}}$ instead. Also, if $\boldsymbol{c} = \boldsymbol{0}$, we omit it in the subscript. We finally define $D_{\Lambda,[\alpha,\alpha']}$ for $0 \leq \alpha < \alpha'$, as the set of discrete Gaussian distributions $D_{\Lambda,\boldsymbol{r}}$ with $\alpha < r_i \leq \alpha'$ for all i. The following states that it is possible to efficiently sample from $D_{\Lambda,\sigma,\boldsymbol{c}}$ when σ is sufficiently large.

Lemma 2 (Adapted from [19, Th. 4.1]). *There exists a probabilistic polynomial-time algorithm that, given a basis $(\boldsymbol{b}_i)_i$ of an n-dimensional lattice Λ, a standard deviation $\sigma \geq \max_i \|\boldsymbol{b}_i\| \cdot \omega(\sqrt{\log n})$ (resp. $\sigma \geq \max_i \|\boldsymbol{b}_i\| \cdot \Omega(\sqrt{n})$), and a center $\boldsymbol{c} \in \mathbb{R}^n$, outputs a sample from a distribution that is within negligible (resp. exponentially small) distance to $D_{\Lambda,\sigma,\boldsymbol{c}}$.*

For a lattice Λ and a real $\varepsilon > 0$, the smoothing parameter $\eta_\varepsilon(\Lambda)$ is the smallest $\sigma > 0$ such that $\rho_{1/\sigma}(\Lambda^* \setminus \{0\}) \leq \varepsilon$. By [31, Le. 3.3], we have that $\eta_\varepsilon(\Lambda) \leq O(\sqrt{\log(n/\varepsilon)}) \cdot \lambda_n(\Lambda)$, for any n-dimensional lattice Λ. This gives that smoothing parameter of the module lattice R^d is $\eta_\varepsilon(R^d) \leq \sqrt{n \log(nd/\varepsilon)}$. As $R^* = \frac{1}{n}R$, we also have that $\eta_\varepsilon((R^*)^d) \leq \sqrt{\log(nd/\varepsilon)/n}$.

We now recall some useful properties on discrete Gaussians.

Lemma 3 (Adapted from [32, Cor. 5.3]). *Let Λ be a n-dimensional lattice, $\varepsilon \in (0,1)$ and $\boldsymbol{r} \in \mathbb{R}^n$ with $r_i \geq \eta_\varepsilon(\Lambda)$ for all $i \leq n$. Then we have $Pr_{\boldsymbol{x} \hookleftarrow D_{\Lambda,\boldsymbol{r}}}[\|\boldsymbol{x}\|_\infty \geq (\max_i r_i) \cdot t] \leq 2en \cdot \exp(-\pi t^2)$ 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. $2^{-\Omega(n)}$).*

Lemma 4 ([19, Cor. 2.8]). *Let $\Lambda' \subseteq \Lambda$ be n-dimensional lattices. Then for any $\varepsilon \in (0,1)$, any $\sigma \geq \eta_\varepsilon(\Lambda')$, and any $\boldsymbol{c} \in \mathbb{R}^n$, we have $\Delta(D_{\Lambda,\sigma,\boldsymbol{c}} \bmod \Lambda', U(\Lambda/\Lambda')) \leq 2\varepsilon$.*

Lemma 5 (Adapted from [38, Claim 3.9]). *Let Λ be an n-dimensional lattice, $\boldsymbol{u} \in \mathbb{R}^n$, $\boldsymbol{r} \in \mathbb{R}^n$ with $r_i > 0$ for all i, $\sigma > 0$ and $t_i = t_{i+n/2} = \sqrt{r_i^2 + \sigma^2}$, for all $i \leq n/2$. Assume that $r_i\sigma/t_i \geq \eta_\varepsilon(\Lambda)$ for all i and some $\varepsilon \in (0, 1/2)$. Consider the continuous distribution Y on $\mathbb{R}^n$ obtained by sampling from $D_{\Lambda+\boldsymbol{u},\boldsymbol{r}}$ and then adding a vector taken from ν_σ. Then we have $\Delta(Y, \nu_{\boldsymbol{t}}) \leq 4\varepsilon$.*

The proof follows the same principle as the one of [38, Claim 3.9]. It is given in appendix.

Lemma 6 (Adapted from [34, Th. 3.1]). *Let σ_1, $\sigma_2 > 0$ be positive. Let σ and σ_3 be defined by $\sigma^2 = \sigma_1^2 + \sigma_2^2$ and $\sigma_3^{-2} = \sigma_1^{-2} + \sigma_2^{-2}$. Let Λ_1, Λ_2 be lattices such that $\sigma_1 \geq \eta_\varepsilon(\Lambda_1)$ and $\sigma_3 \geq \eta_\varepsilon(\Lambda_2)$ for some $\varepsilon \in (0, 1/2]$. Consider the following experiment:*

$$\textit{Sample } \boldsymbol{x}_2 \hookleftarrow D_{\Lambda_2,\sigma_2}, \textit{ then sample } \boldsymbol{x}_1 \hookleftarrow \boldsymbol{x}_2 + D_{\Lambda_1 - \boldsymbol{x}_2,\sigma_1}.$$

Then the marginal distribution of $\boldsymbol{x}_1$ is within statistical distance 8ε of $D_{\Lambda_1,\sigma}$.

Lemma 7 ([24, Le. 10]). *Let $\boldsymbol{r} \in \mathbb{R}^n$ with $r_i = r_{i+n/2}$ for all $i \leq n/2$, $\boldsymbol{x} \in K^d$ sampled from $\sigma_H^{-1}(\nu_{\boldsymbol{r},\ldots,\boldsymbol{r}})$, and $\boldsymbol{s} \in K^d$. Then $\sigma_H(\langle \boldsymbol{x}, \boldsymbol{s} \rangle)$ has distribution $\nu_{\boldsymbol{r}'}$ with $r'_i = r'_{i+n/2} = r_i \cdot \sqrt{\sum_{k=1}^d |\sigma_i(s_k)|^2}$, for all $i \leq n/2$.*

2.3 Computational problems

Let $\gamma \geq 1$ be a function of the dimension n. The approximate *Shortest Independent Vectors Problem* SIVP$_\gamma$ is as follows: Given a basis $(\boldsymbol{b}_i)_i$ of an n-dimensional lattice Λ, find n linearly independent vectors $(\boldsymbol{s}_i)_i$ in Λ such that $\max_i \|\boldsymbol{s}_i\| \leq \gamma \cdot \lambda_n(\Lambda)$. We let Id-SIVP (resp. Mod-SIVP$^{(n,d)}$) denote the restriction of SIVP to ideal (resp. (n,d)-module) lattices.

The MLWE problem generalizes both LWE and RLWE. It was introduced in [11] and a reduction from approximate Mod-SIVP was given in [24]. Let R_q and R_q^* respectively denote R/qR and R^*/qR^*. Let ψ be some distribution on $\mathbb{T}_{R^*} := K_\mathbb{R}/R^*$ and $\boldsymbol{s} \in (R_q^*)^d$ be an arbitrary vector. We define $A_{\boldsymbol{s},\psi}^{(q)}$ as the distribution on $R_q^d \times \mathbb{T}_{R^*}$ obtained by sampling a vector $\boldsymbol{a}$ uniformly in R_q^d, and $e \in \mathbb{T}_{R^*}$ from distribution ψ, and returning $(\boldsymbol{a}, \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s} \rangle + e)$.

Definition 1. *Let $n, d \geq 1$ and $q \geq 2$ be integers, and Υ be a distribution over a family of distributions over $K_{\mathbb{R}}$. The* decision *version of the* Module Learning With Error problem $\mathrm{MLWE}_{q,\Upsilon}^{(n,d)}$ *is as follows: Let $\boldsymbol{s} \in (R_q^*)^d$ be sampled uniformly and ψ be sampled from Υ; The goal is to distinguish between arbitrarily many independent samples from $A_{\boldsymbol{s},\psi}^{(q)}$ and the same number of independent samples from $U(R_q^d \times \mathbb{T}_{R^*})$.*

When $n = 1$, we recover the LWE problem, whereas RLWE corresponds to setting $d = 1$. Note that the noise distribution is itself chosen randomly (from distribution Υ). This randomization of the noise distribution was introduced in [27], to obtain a reduction from a worst-case lattice problem to a fully average-case variant of RLWE (without it, the reduction from [27] would be to a variant of RLWE with noise distribution arbitrarily chosen in a set of distributions). In the case of LWE, the sampling of the noise can be avoided using [38, Le. 3.7], and in the case of RLWE with a bounded number of samples (which is the case in most applications), it can be avoided using [28, Se. 5.2].

The following result from [24] generalizes both that of [38] on LWE, and that of [27] on RLWE.

Theorem 1 (Adapted from [24, Th. 9]). *Let $d, n \geq 1$, $\varepsilon = (nd)^{-\omega(1)}$ (resp. $\varepsilon = 2^{-\Omega(nd)}$), $\alpha \in (0, 1)$ and q prime such that $q = 1 \bmod 2n$ and $\alpha q > \omega(\sqrt{d \log n})$ (resp. $\alpha q > \Omega(\sqrt{dn})$). There exists a quantum reduction from solving* $\text{Mod-SIVP}_{\gamma}^{(n,d)}$ *in polynomial-time with non-negligible probability (resp. sub-exponential time with non-exponentially small probability) to solving* $\mathrm{MLWE}_{q,\Upsilon_\alpha}^{(n,d)}$ *in polynomial-time with non-negligible probability (resp. sub-exponential time with non-exponentially small probability), with $\gamma = \frac{\sqrt{n}d}{\alpha} \sqrt{\log(nd)} \cdot \omega(\sqrt{\log n})$ (resp. $\frac{1}{\alpha} \cdot \Omega((nd)^{3/2})$).*

In [33], Peikert proposed a dequantumized variant of Theorem 1, from standard lattice problems such as GapSVP to decision LWE (i.e., $\mathrm{MLWE}^{(1,n)}$), when $q = 2^{\Omega(\frac{n \log\log n}{\log n})}$ is a product of distinct primes that are polynomial in n. This was later improved in [30] to handle for powers of small primes.

Assume that $\alpha \geq \sqrt{\log(nd/\varepsilon)}/n^{3/4}$. Then by [31, Le. 4.1] and the bound on $\eta_\varepsilon((R^*)^d)$ from Subsection 2.2, we have that $\Delta(A_{\boldsymbol{s},\psi}^{(q)}, U(R_q^d \times \mathbb{T}_{R^*})) \leq \varepsilon/2$ with overwhelming probability with respect to $\psi \hookleftarrow \Upsilon_\alpha$ and $\boldsymbol{s} \hookleftarrow U(R_q^d)$. This implies that for $\alpha \geq \omega(\sqrt{\log(nd)})/n^{3/4}$ (resp. $\alpha \geq \Omega(d^{1/2}/n^{1/4})$, solving $\mathrm{MLWE}_{q,\Upsilon_\alpha}^{(n,d)}$ in polynomial time (resp. sub-exponential time) in nd is vacuously hard.

3 A modulus-switching self-reduction for MLWE

The aim of the present section is to prove the following result.

Theorem 2. *Let $d, n \geq 1$, $p, q \in [2, 2^{(nd)^{O(1)}}]$ and $\alpha, \beta \in (0, 1)$ such that $\beta \geq \alpha \cdot \max(1, \frac{q}{p}) \cdot n^{3/4} d^{1/2} \cdot \omega(\log^2 nd)$ (resp. $\beta \geq \alpha \cdot \max(1, \frac{q}{p}) \cdot \Omega(n^{11/4} d^{5/2})$) and $\alpha q \geq \omega(\sqrt{\log(nd)/n})$ (resp. $\Omega(\sqrt{d})$). There exists a probabilistic polynomial-time reduction from solving* $\mathrm{MLWE}_{q,\Upsilon_\alpha}^{(n,d)}$ *in polynomial time with non-negligible probability (resp. sub-exponential time with non-exponentially small probability) to solving* $\mathrm{MLWE}_{p,\Upsilon_\beta}^{(n,d)}$ *in polynomial time with non-negligible probability (resp. sub-exponential time with non-exponentially small probability).*

Note that the condition on αq from Theorem 2 is always much weaker than the one from Theorem 1. Combined with Theorem 1, Theorem 2 provides a polynomial-time quantum reduction from (worst-case) Mod-SIVP to MLWE with a modulus p of arbitrary arithmetic form. As MLWE is a generalization of both LWE and RLWE, we recall that this theorem also provides a reduction from SIVP_γ (resp. Id-SIVP_γ) to $\mathrm{LWE}_{p,\Upsilon_\beta}$ (resp. $\mathrm{RLWE}_{p,\Upsilon_\beta}$), for a modulus p of arbitrary arithmetic shape. For instance, SIVP_γ reduces to $\mathrm{LWE}_{p,\Upsilon_\beta}$ as soon as $\beta p = \widetilde{\Omega}(n)$ and $\gamma = \widetilde{\Omega}(n^{3/2}/\beta)$ (by choosing $q \approx n/\beta$ prime and $\alpha \approx \beta/n^{1/2}$).

Note that in the case of LWE, Regev [38] showed a hardness result for a fixed error distribution ν_α (as opposed to a randomly chosen error distribution). Therefore, we may start the modulus-switching reduction from instances of $\mathrm{LWE}_{q,\nu_\alpha}$, and allows one to save a $\log d$ factor in the requirement on β/α from Theorem 2.

The proof of Theorem 2 proceeds by a sequence of reductions:

$$\mathrm{MLWE}_{q,\Upsilon_\alpha} \xrightarrow{\text{Section 3.1}} \text{HNF-MLWE}_{q,D_{\frac{1}{q}R^*,[\alpha,\alpha']}} \xrightarrow{\text{Section 3.2}} \mathrm{MLWE}_{p,\Psi_{\leq\beta}} \xrightarrow{\text{Lemma 13}} \mathrm{MLWE}_{p,\Upsilon_\beta}$$

We first reduce $\mathrm{MLWE}_{q,\Upsilon_\alpha}$ to the HNF version of $\mathrm{MLWE}_{q,D_{(1/q)R^*,[\alpha,\alpha']}}$ (i.e., with a small secret $\boldsymbol{s}$), where $\alpha' \approx \alpha n^{1/4}$. Then we reduce HNF-$\mathrm{MLWE}_{q,D_{(1/q)R^*,[\alpha,\alpha']}}$ to $\mathrm{MLWE}_{p,\Psi_{\leq\beta}}$, by switching the modulus and handling the Right Hand Sides of the MLWE samples so that the error term comes from some distribution in the set $\Psi_{\leq\beta}$. Finally, we re-randomize the noise distribution, thus providing a reduction from $\mathrm{MLWE}_{p,\Psi_{\leq\beta}}$ to $\mathrm{MLWE}_{p,\Upsilon_\beta}$.

3.1 Reducing $\mathbf{MLWE}_{q,\Upsilon_\alpha}$ to HNF-$\mathbf{MLWE}_{q,D_{\frac{1}{q}R^*,[\alpha,\alpha']}}$

We consider a sample ϕ from distribution Υ_α. Recall that if ϕ is sampled from Υ_α, then $\phi = \nu_{\boldsymbol{r}}$ with $r_i = r_{i+n/2} = \alpha\sqrt{1+\sqrt{n}x_i}$ and x_i sampled from $\Gamma(2,1)$, for all $i \leq n/2$. By definition of $\Gamma(2,1)$, we have that, $\Pr_{x \hookleftarrow \Gamma(2,1)}[x \leq y] = 1 - (1+y)e^{-y}$. We derive that $x \leq \omega(\log nd)$ (resp. $x \leq \Omega(nd)$) with probability negligibly (resp. exponentially) close to 1. As a consequence, with the same probability we have that $\alpha < r_i \leq \alpha' = \alpha \cdot n^{1/4}\omega(\log nd)$ (resp. $\alpha' = \alpha \cdot \Omega(n^{5/4}d)$) for all i. Therefore, $\mathrm{MLWE}_{q,\Psi_{[\alpha,\alpha']}}$ is as least as hard as $\mathrm{MLWE}_{q,\Upsilon_\alpha}$.

We now discretize the noise distribution.

Lemma 8 (Adapted from [20, Le. 2]). *For any n, d, q, $\varepsilon \in (0,1)$, $\boldsymbol{r} \in \mathbb{R}^n$ and α satisfying $r_i > \alpha$ for all i and $\alpha q \geq \eta_\varepsilon(R^*)$, $\mathrm{MLWE}_{q,D_{(1/q)R^*,\sqrt{2}\boldsymbol{r}}}$ reduces (in polynomial time) to $\mathrm{MLWE}_{q,\nu_{\boldsymbol{r}}}$.*

The proof is following the same principle as the proof of [20, Le. 2]. It is given in appendix.

The following lemma allows us to reduce the $\mathrm{MLWE}_{q,D_{(1/q)R^*,\sqrt{2}\boldsymbol{r}}}$ problem to a variant in which the secret is chosen from $D_{(R^*)^d,\sqrt{2}q\boldsymbol{r}}$. We call this new problem the Hermite Normal Form (HNF) of MLWE.

Lemma 9 (Adapted from [7, Le. 2]). *There exists a deterministic polynomial time transformation that, for arbitrary $\boldsymbol{s} \in (R_q^*)^d$ and error distribution $D_{(1/q)R^*,\boldsymbol{r}}$, maps $A^{(q)}_{\boldsymbol{s},D_{(1/q)R^*,\boldsymbol{r}}}$ to $A^{(q)}_{\overline{\boldsymbol{x}},D_{(1/q)R^*,\boldsymbol{r}}}$ where $\overline{\boldsymbol{x}} \hookleftarrow D_{(R^*)^d,q\boldsymbol{r}}$, and maps $U((R_q)^d \times \mathbb{T}_{R^*})$ to itself.*

The proof is following the same principle as the proof of [7, Le. 2]. It is given in appendix.

This completes the reduction from $\mathrm{MLWE}_{q,\Upsilon_\alpha}$ to the HNF of $\mathrm{MLWE}_{q,D_{(1/q)R^*,[\alpha,\alpha']}}$.

3.2 Reducing HNF-$\mathbf{MLWE}_{q,D_{\frac{1}{q}R^*,[\alpha,\alpha']}}$ to $\mathbf{MLWE}_{p,\Psi_{\leq\beta}}$

This is the main component of the proof of Theorem 2. We first show how to transform a sample from $A^{(q)}_{\boldsymbol{s},D_{(1/q)R^*,[\alpha,\alpha']}}$ to a sample of $A^{(p)}_{\boldsymbol{s},\Psi_{\leq\beta}}$.

Handling the first component of the MLWE sample. By definition, the element $\boldsymbol{a}$ is chosen uniformly at random in $(R_q)^d$. The following randomized function allows us to map the first part of the MLWE distribution $\boldsymbol{a}$ to a uniform element in $(R_p)^d$.

Lemma 10. *Let $n, d \geq 1$, $p, q \geq 2$ be integers, $\varepsilon \in (0,1)$ and $\sigma \geq \max\left(1, \sqrt{\frac{p^2}{q^2-1}}\right) \cdot \eta_\varepsilon(R^d)$. We define the following randomized function $f : (R_q)^d \to (R_p)^d$:*

$$f(\boldsymbol{a}) = \frac{p}{q}\boldsymbol{a} + D_{R^d - \frac{p}{q}\boldsymbol{a},\sigma}.$$

Then, if $\boldsymbol{a}$ is sampled uniformly in $(R_q)^d$ then the distribution of $f(\boldsymbol{a})$ is within statistical distance 12ε to the uniform distribution in $(R_p)^d$.

Proof. We take a parameter $\sigma_2' = \eta_\varepsilon(qR^d)$, and apply Lemma 4 with $\Lambda = R^d$ and $\Lambda' = qR^d$. We obtain that the statistical distance between the uniform distribution on $(R_q)^d$ and the distribution $(D_{R^d,\sigma_2'} \bmod qR^d)$ is at most 2ε.

As a consequence, for any integer $p \geq 2$, the distribution $\frac{p}{q}U((R_q)^d)$ is within statistical distance 2ε to the distribution $\frac{p}{q}(D_{R^d,\sigma_2'} \bmod qR^d)$. The latter is the same distribution as $(\frac{p}{q}D_{R^d,\sigma_2'} \bmod pR^d)$. We also have that $\frac{p}{q}D_{R^d,\sigma_2'} = D_{\frac{p}{q}R^d,\sigma_2}$, with $\sigma_2 = \frac{p}{q}\sigma_2'$.

Let $\boldsymbol{a}$ be uniformly distributed in $(R_q)^d$, so far we have shown that $\frac{p}{q}\boldsymbol{a}$ is within statistical distance at most 2ε from $(D_{\frac{p}{q}R^d,\sigma_2} \bmod pR^d)$. Now, we use Lemma 6, and write:

$$\Lambda_1 = R^d, \quad \Lambda_2 = \frac{p}{q}R^d, \quad \sigma_2 = p \cdot \eta_\varepsilon(R^d), \quad \text{and} \quad \sigma_1 = \sigma \geq \max\left(1, \sqrt{\frac{p^2}{q^2-1}}\right) \cdot \eta_\varepsilon(R^d).$$

We have:

- $\sigma'^2 := \sigma_1^2 + \sigma_2^2 \geq \max\left(p^2+1, \frac{p^2q^2}{q^2-1}\right) \cdot \eta_\varepsilon^2(R^d)$,
- $\sigma_3^{-2} := \sigma_1^{-2} + \sigma_2^{-2}$ and $\sigma_1 \geq \sqrt{\frac{p^2}{q^2-1}} \cdot \eta_\varepsilon(R^d)$, thus $\sigma_3 \geq \frac{p}{q} \cdot \eta_\varepsilon(R^d)$.

Overall, we have $\sigma_1 \geq \eta_\varepsilon(R^d)$ and $\sigma_3 \geq \frac{p}{q} \cdot \eta_\varepsilon(R^d)$, and therefore the assumptions of Lemma 6 hold. Applying the lemma gives that the residual distribution of $\boldsymbol{x}_1$ after the following experiment is within statistical distance 8ε of $D_{R^d,\sigma'}$:

$$\text{Sample } \boldsymbol{x}_2 \hookleftarrow D_{\frac{p}{q}R^d,\sigma_2}, \quad \text{then sample } \boldsymbol{x}_1 \hookleftarrow \boldsymbol{x}_2 + D_{R^d - \boldsymbol{x}_2,\sigma}.$$

As a consequence, the vector $f(\boldsymbol{a})$ is within statistical distance 10ε of $(D_{R^d,\sigma'} \bmod pR^d)$.

Finally, as $\sigma' \geq p \cdot \eta_\varepsilon(R^d) = \eta_\varepsilon(pR^d)$, Lemma 4 implies that the statistical distance between $(D_{R^d,\sigma'} \bmod pR^d)$ and the distribution $U((R_p)^d)$ is at most 2ε. By combining this with the above analysis of the distribution of $f(\boldsymbol{a})$, we obtain the statistical distance between the distribution of $f(\boldsymbol{a})$ and the uniform distribution on $(R_p)^d$ is at most 12ε. □

For the following, we set $\sigma = \sqrt{2}\max\left(1, \sqrt{\frac{p^2}{q^2-1}}\right) \cdot \eta_\varepsilon(R^d)$.

Handling the second component of the MLWE sample. Let $b = \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s}\rangle + e$, with $\boldsymbol{a}$ sampled uniformly in $(R_q)^d$, $\boldsymbol{s}$ fixed and e sampled from $D_{(1/q)R^*,\boldsymbol{r}}$, where $\alpha < r_i = r_{i+n/2} \leq \alpha'$, for all $i \leq n/2$. Let $\boldsymbol{a}' = f(\boldsymbol{a})$ be the element of $(R_p)^d$ obtained from $\boldsymbol{a}$ as described above. Then:

$$b = \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s}\rangle + e = \frac{1}{p}\langle \boldsymbol{a}', \boldsymbol{s}\rangle + \frac{1}{p}\langle \frac{p}{q}\boldsymbol{a} - \boldsymbol{a}', \boldsymbol{s}\rangle + e.$$

We consider the new error term $\frac{1}{p}\langle \frac{p}{q}\boldsymbol{a} - \boldsymbol{a}', \boldsymbol{s}\rangle + e$. We aim at transforming it into a continuous Gaussian distribution with parameter related to β. As it is discrete, we add to it continuous Gaussians to smooth it. In fact, we are to consider both components $\frac{1}{p}\langle \frac{p}{q}\boldsymbol{a} - \boldsymbol{a}', \boldsymbol{s}\rangle$ and e independently. To smooth the first component, we add to it a sample e_d from $\nu_{\sigma \cdot s_{max}}$, where $s_{\max} = \sqrt{d} \cdot \alpha' q \cdot \omega(\sqrt{\log nd})$ (resp. $s_{\max} = \alpha' q \cdot \Omega(\sqrt{n}d)$). We will consider the second component directly in the description of the reduction from HNF-MLWE$_{q,D_{(1/q)R^*,[\alpha,\alpha']}}$ to MLWE$_{p,\Psi_{\leq\beta}}$.

First, we know that $\|\boldsymbol{s}\|_{2,\infty} \leq \sqrt{d}\|\boldsymbol{s}\|_\infty$. Let $\varepsilon = (nd)^{-\omega(1)}$ (resp. $\varepsilon = 2^{-\Omega(nd)}$), by assumption, we have that $\alpha q \geq \eta_\varepsilon(R^*)$. By Lemma 3 we have that $\|\boldsymbol{s}\|_\infty \leq \alpha' q \cdot \omega(\sqrt{\log nd})$ (resp. $\|\boldsymbol{s}\|_\infty \leq \alpha' q \cdot \Omega(\sqrt{nd})$) with probability $\geq 1 - \varepsilon$, for $\boldsymbol{s}$ sampled from any $D_{\Lambda,\boldsymbol{r}}$ such that $r_i \in (\alpha, \alpha']$ for all i. As a consequence, with the same probability we have that $\|\boldsymbol{s}\|_{2,\infty} \leq s_{\max}$.

Lemma 11. *Let $S > 0$ and $\boldsymbol{s} \in K^d$ with $\|\boldsymbol{s}\|_{2,\infty} < S$. Let $\boldsymbol{d}$ be distributed as $D_{R^d - \boldsymbol{a},\sigma}$ for some arbitrary $\boldsymbol{a}$ and $\sigma \geq \sqrt{2}\eta_\varepsilon(R^d)$ and e be distributed as ν_τ for some $\tau \geq \sigma \cdot S$. Then the distribution of $\langle \boldsymbol{d}, \boldsymbol{s}\rangle + e$*

is within statistical distance 4ε of the elliptical Gaussian distribution $\nu_{\boldsymbol{t}}$ over K, where $t_i^2 = t_{i+n/2}^2 = \sigma^2 \sum_{k=1}^d |\sigma_i(s_k)|^2 + \tau^2$, for all $i \le n/2$.[6]

Proof. By Lemma 7, we have that e is following the same distribution as $\langle \boldsymbol{e}_s, \boldsymbol{s} \rangle$ with $\boldsymbol{e}_s$ distributed from $\nu_{\mathbf{r}',\ldots,\mathbf{r}'}$ and $r'_i = r'_{i+n/2} = \tau / \sqrt{\sum_{k=1}^d |\sigma_i(s_k)|^2}$ for $i \le n/2$.

As a consequence, we have that $\langle \boldsymbol{d}, \boldsymbol{s} \rangle + e$ is following the same distribution as $\langle \boldsymbol{d} + \boldsymbol{e}_s, \boldsymbol{s} \rangle$. We write $\boldsymbol{e}_s = \boldsymbol{e}_1 + \boldsymbol{e}_2$ with $\boldsymbol{e}_1$ distributed from $\nu_{\tau/S}$ and $\boldsymbol{e}_2$ distributed from $\nu_{(\sqrt{(r'_i)^2 - (\tau/S)^2})_i}$. We now use Lemma 5: As $\sigma \ge \sqrt{2}\eta_\varepsilon(R^d)$ and $\tau \ge \sqrt{2} S \cdot \eta_\varepsilon(R^d)$, we have that $\boldsymbol{d} + \boldsymbol{e}_1$ is within statistical distance 4ε from $\nu_{\sqrt{\sigma^2 + (\tau/S)^2}}$. Now, the quantity $\boldsymbol{d} + \boldsymbol{e}_s$ can be interpreted as the sum of two continuous Gaussians: It is within statistical distance 4ε from $\nu_{(\sqrt{\sigma^2 + (r'_i)^2})_i}$.

We use Lemma 7 once more. We obtain that $\langle \boldsymbol{d}, \boldsymbol{s} \rangle + e$ is within statistical distance 4ε from $\nu_{\boldsymbol{t}}$ with $t_i^2 = t_{i+n/2}^2 = \sigma^2 \sum_{k=1}^d |\sigma_i(s_k)|^2 + \tau^2$, for all $i \le n/2$. □

Reduction from the HNF of $\mathrm{MLWE}_{q, D_{(1/q)R^*, [\alpha, \alpha']}}$ to $\mathrm{MLWE}_{p, \Psi_{\le \beta}}$. Assume that we have a polynomial time (resp. sub-exponential time) oracle that solves $\mathrm{MLWE}_{p, \Psi_{\le \beta}}$ with probability $(nd)^{-O(1)}$ (resp. $2^{-o(nd)}$). Our inputs are m samples from either $U(R_q^d \times \mathbb{T}_{R^*})$ or from $A^{(q)}_{\boldsymbol{s}, D_{(1/q)R^*, \mathbf{r}}}$, where the secret $\boldsymbol{s}$ is distributed from $D_{(R^*)^d, q\mathbf{r}}$ and where $\alpha < r_i = r_{i+n/2} \le \alpha'$ for all i. The reduction $\mathcal{R}$ is as follows:

- Sample $\boldsymbol{s}_1$ uniformly in $(R_p^*)^d$.
- For each sample $(\boldsymbol{a}, b)$, create the sample $(\boldsymbol{a}', b')$ as follows:
    - Let e' be a sample of $\nu_{\alpha'}$ and let e_d be a sample of $\nu_{\sigma \cdot s_{\max}}$.
    - Let $\boldsymbol{a}' = f(\boldsymbol{a})$ as defined in Lemma 10 and $b' = b + \frac{1}{p}\langle \boldsymbol{a}', \boldsymbol{s}_1 \rangle + \frac{1}{p} e_d + e'$.
    - Return the new sample $(\boldsymbol{a}', b')$.
- Call the $\mathrm{MLWE}_{p, \Psi_{\le \beta}}$ oracle on these samples, and return its answer.

Lemma 12. *Let $\varepsilon = (nd)^{-\omega(1)}$ (resp. $\varepsilon = 2^{-\Omega(nd)}$). The following holds with probability $\ge 1 - \varepsilon$. If the samples $(\boldsymbol{a}, b)$ are from $A^{(q)}_{\boldsymbol{s}, D_{(1/q)R^*, r}}$, then the derived samples $(\boldsymbol{a}', b')$ follow a distribution that is within statistical distance 20ε of $A^{(p)}_{\boldsymbol{s}', \nu_{\boldsymbol{t}''}}$, where $\boldsymbol{s}'$ is uniform in $(R_p^*)^d$ and $0 < t''_i \le \beta$ for all i.*

Proof. We assume that $\|\boldsymbol{s}\|_{2,\infty} \le s_{max}$, which holds with probability $\ge 1 - \varepsilon$.

We showed that $\boldsymbol{a}'$ is within statistical distance 12ε from the uniform distribution in $(R_q)^d$. Also, we have that $b' = \frac{1}{p}\langle \boldsymbol{a}', \boldsymbol{s} + \boldsymbol{s}_1 \rangle + \frac{1}{p}\langle \frac{p}{q}\boldsymbol{a} - \boldsymbol{a}', \boldsymbol{s} \rangle + \frac{1}{p} e_d + e + e'$.

First, we have that $\boldsymbol{s}' = \boldsymbol{s} + \boldsymbol{s}_1$, where $\boldsymbol{s}_1$ is uniform in $(R_p^*)^d$ and independent from $\boldsymbol{s}$. This ensures that $\boldsymbol{s}' \bmod p$ is uniform in $(R_p^*)^d$.

We now study the component $\langle \frac{p}{q}\boldsymbol{a} - \boldsymbol{a}', \boldsymbol{s} \rangle + e_d$. By applying Lemma 11, we have that it is within statistical distance 4ε from $\nu_{\boldsymbol{t}}$ with $t_i^2 = t_{i+n/2}^2 = \sigma^2 \left(\sum_{k=1}^d |\sigma_i(s_k)|^2 + s_{\max}^2 \right)$. As a consequence, the quantity $\frac{1}{p}(\langle \frac{p}{q}\boldsymbol{a} - \boldsymbol{a}', \boldsymbol{s} \rangle + e_d)$ is within statistical distance 4ε from $\nu_{\frac{1}{p}\boldsymbol{t}}$.

By hypothesis $\alpha' q \ge r_i q > \alpha q \ge \sqrt{2}\eta_\varepsilon(R^*)$ holds for all i. Thus, by Lemma 5, we have that $e'' = e + e'$ is within statistical distance 4ε from $\nu_{\boldsymbol{t}'}$ with $(t'_i)^2 = (t'_{i+n/2})^2 = r_i^2 + (\alpha')^2$. Finally, the error component $\frac{1}{p}(\langle \frac{p}{q}\boldsymbol{a} - \boldsymbol{a}', \boldsymbol{s} \rangle + e_d) + e''$ is within statistical distance 8ε from $\nu_{\boldsymbol{t}''}$ with $(t''_i)^2 = (t''_{i+n/2})^2 = r_i^2 + (\alpha')^2 + \frac{\sigma^2}{p^2}(\sum_{k=1}^d |\sigma_i(s_k)|^2 + s_{\max}^2)$.

We now bound the t''_i's. As $r_i \le \alpha'$ holds for all i, and using the fact that $\|\boldsymbol{s}\|_{2,\infty} \le s_{max}$, we have:

$$t''_i = t''_{i+n/2} \le \sqrt{2}\alpha \cdot n^{1/4} \omega(\log nd) \cdot \sqrt{1 + \frac{q^2}{p^2}\sigma^2 d \cdot \omega(\log nd)} \le \beta \quad \text{for all } i$$

[6] To be rigorous, we actually let $\boldsymbol{d}$ be distributed as $\sigma_H^{-1}(D_{R^d - \boldsymbol{a}, \sigma})$, and consider the distribution of $\sigma_H(\langle \boldsymbol{d}, \boldsymbol{s} \rangle)$. For simplicity, we identify elements of K^d with their σ_H embeddings.

(resp. $t''_i = t''_{i+n/2} \leq \sqrt{2}\alpha \cdot \Omega(n^{5/4}d) \cdot \sqrt{1 + \frac{q^2}{p^2}\sigma^2 \Omega(nd^2)} \leq \beta$ for all i). □

Conversely, if the samples $(\boldsymbol{a}, b)$ are from the uniform distribution over $R_q^d \times \mathbb{T}_{R^*}$, then the resulting samples $(\boldsymbol{a}', b')$ are within statistical distance 12ε from the uniform distribution over $R_p^d \times \mathbb{T}_{R^*}$.

The above arguments imply the correctness of reduction $\mathcal{R}$. Let $\varepsilon = (nd)^{-\omega(1)}$ (resp. $\varepsilon = 2^{-\Omega(nd)}$), then with probability of success $(nd)^{-O(1)}$ (resp. $2^{-o(nd)}$), it allows us to solve the instance of the HNF version of MLWE. There are two possible kinds of inputs. In the first case, it is given as inputs uniformly distributed samples, whereas in the second case, it is given as inputs samples from a valid $A^{(q)}_{\boldsymbol{s},D_{(1/q)R^*,\mathbf{r}}}$. With probability $\geq 1 - \varepsilon$, the transformation on the samples achieves the following: The uniform samples remain (essentially) uniform and the samples from $A^{(q)}_{\boldsymbol{s},D_{(1/q)R^*,\mathbf{r}}}$ become samples within statistical distance 20ε of a valid distribution $A^{(p)}_{\boldsymbol{s}',\nu_{\boldsymbol{t}''}}$ for $\mathrm{MLWE}_{p,\Psi_{\leq\beta}}$. As a consequence, the $\mathrm{MLWE}_{p,\Psi_{\leq\beta}}$ oracle is given a valid input, and succeeds (with good advantage) in correctly guessing which distribution it is given. The advantage of reduction $\mathcal{R}$ in correctly solving the HNF of $\mathrm{MLWE}_{q,D_{(1/q)R^*,[\alpha,\alpha']}}$ is not less than 20ε of the advantage of the $\mathrm{MLWE}_{p,\Psi_{\leq\beta}}$ oracle.

3.3 Reducing $\mathrm{MLWE}_{p,\Psi_{\leq\beta}}$ to $\mathrm{MLWE}_{p,\Upsilon_\beta}$

This reduction is the last component of the proof of Theorem 2. The goal is to re-randomize the error distribution of MLWE. The proof is adapted from the proof of [28, Le. 5.11].

Lemma 13. *Let $p \geq 2$ be an integer and $\beta \in (0,1)$. There exists a randomized polynomial-time reduction from solving $\mathrm{MLWE}_{p,\Psi_{\leq\beta}}$ in polynomial-time (resp. sub-exponential time) with non-negligible (resp. non-exponentially small) probability to solving $\mathrm{MLWE}_{p,\Upsilon_\beta}$ in polynomial-time (resp. sub-exponential time) with non-negligible (resp. non-exponentially small) probability.*

Proof. Let $(\boldsymbol{a}, b = \frac{1}{p}\langle \boldsymbol{a}, \boldsymbol{s}\rangle + e)$ be a sample from $A^{(p)}_{\boldsymbol{s},\nu_{\boldsymbol{t}}}$ with $0 < t_i \leq \beta$ for all i and $\boldsymbol{s} \hookleftarrow U((R_q^*)^d)$. Let $x'_1, \ldots, x'_{n/2}$ be independent samples from $\Gamma(2,1)$. We perform the following transformation:

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

where e' is sampled from $\nu_{\boldsymbol{r}}$, with $\boldsymbol{r}$ defined by $r_i^2 = r_{i+n/2}^2 = \beta^2\sqrt{n}x'_i$.

This transformation maps the uniform distribution over $(R_p)^d \times \mathbb{T}_{R^*}$ to itself. On the other hand, it maps $A^{(p)}_{\boldsymbol{s},\nu_{\boldsymbol{t}}}$ to $A^{(p)}_{\boldsymbol{s},\nu_{\boldsymbol{r}'}}$, with $r'_i = r'_{i+n/2} = \sqrt{t_i^2 + \beta^2\sqrt{n}x'_i}$, for all $1 \leq i \leq n/2$.

Let S denote the set of ψ's for which the oracle distinguishes with non-negligible (resp. non-exponentially small) probability between the uniform distribution over $(R_p)^d \times \mathbb{T}_{R^*}$ and the distribution $A_{\boldsymbol{s},\psi}$. By assumption, the measure of S under Υ_β is non-negligible (resp. non-exponentially small). Lemma 1 implies that $\nu_{\boldsymbol{r}'} \in S$ with non-negligible (resp. non-exponentially small) probability. The result follows. □

This lemma completes the proof of Theorem 2.

4 Cryptographic applications

As mentioned in the introduction, the proven hardness of RLWE for a modulus q so that R_q is isomorphic to $(\mathbb{F}_{q^{n/2}})^2$ allows us to adapt to the RLWE setting a number of LWE-based schemes, and therefore to significantly improve their performance. In this section, we only consider the application to fast Identity-Based Encryption (IBE). We describe a fast RLWE-based authentication scheme in Appendix B. We will only describe the schemes, without giving their security proofs, as these can be obtained by adapting the existing ones from [1] and [21] in a direct manner. Other LWE-based schemes that can be adapted to RLWE for such a choice of modulus q include [15,40,4,25,36].

Efficient computations in R_q. Let us comment on the efficiency of the specific algebraic setting. Let n be a power of 2, and q be prime such that $q = 3 \bmod 8$. This implies that $x^n + 1$ has exactly two irreducible factors Φ_1 and Φ_2 modulo q, each of degree $n/2$. We interpret the rings K, R and R_q as the polynomial rings $\mathbb{Q}[x]/(x^n+1)$, $\mathbb{Z}[x]/(x^n+1)$ and $\mathbb{Z}_q[x]/(x^n+1)$ respectively. The vector $\sigma_H(y)$ (for $y \in K$) and the vector of the coefficients of y when interpreted as a polynomial are related by a similarity of center 0 and factor $\sqrt{n}$. That similarity is closely related to the $2n$-dimensional complex Fourier transform. This interpretation allows for an efficient conversion of a sample from a noise distribution $\sigma_H^{-1}(\nu_{\boldsymbol{r}})$ to the polynomial setting (see [41, Se. 2] for more details). Once the element has been mapped to $\mathbb{Q}[x]/(x^n + 1)$, it can be multiplied by q and rounded to a closest element of $\mathbb{Z}[x]/(x^n+1)$ (by rounding each coefficient to a nearest integer). The resulting sample is expected to have norm $\approx \alpha q n^{1/4}$. We let $\overline{\Upsilon}_{\alpha,q}$ denote this overall process. Independently, arithmetic in the ring R_q can be efficiently implemented by using standard quasi-linear time polynomial arithmetic [16]. A practical solution could be to choose a prime q' such that $q' = 1 \bmod 2n$ and $q' > 2nq^2$, and to perform multiplications in R_q by first using a fast discrete Fourier transform to multiply in $R_{q'} \simeq (\mathbb{F}_{q'})^n$ and then reduce the coefficients modulo q. Overall, if q is also set such that $q \leq n^{O(1)}$, then sampling from the distribution on noise distributions, sampling from the resulting noise distribution and adding/multiplying elements of R_q can all be performed in time quasi-linear in n.

A fast identity-based encryption scheme. The IBE scheme from Figure 1 is an adaptation of the IBE scheme of Agrawal et al [1] to the RLWE setting. As in [1], the scheme can be adapted to provide a selectively secure hierarchical IBE (with a small number of levels). By combining it with [9] along with a fast one-time signature such as [26], it is possible to derive an IND-CCA (hierarchical identity-based) encryption scheme with quasi-optimal performance and security relying (in the standard model) on the presumed quantum worst-case hardness of approximate Id-SIVP.

- **Master key generation.** Use the algorithm from [42, Se. 3] to obtain $(T, \boldsymbol{a}_0) \in R^{m\times m} \times R_q^m$. Sample $\boldsymbol{a}_1, \boldsymbol{b} \hookleftarrow U(R_q^m)$, and $u \hookleftarrow U(R_q)$. The public parameters are $MPK = (\boldsymbol{a}_0, \boldsymbol{a}_1, \boldsymbol{b}, u)$. The master secret key is $msk = T$.
- **U_{id}'s public key generation.** Use MPK to compute $\boldsymbol{a}_{id} \in R_q^{2m}$ as the concatenation of $\boldsymbol{a}_0$ and $\boldsymbol{a}_1 + \mathcal{H}(id) \cdot \boldsymbol{b}$. The public-key of U_{id} is $pk_{id} = (\boldsymbol{a}_{id}, u)$.
- **U_{id}'s secret key derivation.** Use the key delegation mechanism from [14] to derive from msk a matrix $T_{id} \in R^{2m\times 2m}$ with small entries, full K-rank, and such that $T_{id}\boldsymbol{a}_{id} = \boldsymbol{0} \bmod q$.
- **Encrypting $M \in \mathcal{P}$ for U_{id}.** Use MPK and id to derive $pk_{id} = (\boldsymbol{a}_{id}, u)$. Sample the standard deviations of the noise distribution $\overline{\Upsilon}_{\alpha,q}$. Sample $e_1, \ldots, e_{2m+1} \in R$ from the resulting noise distribution. Let $\boldsymbol{e}$ denote $(e_i)_{i\leq 2m}$. Sample s uniformly in R_q. Return ciphertext $(\boldsymbol{c}_1, c_2) = (s \cdot \boldsymbol{a}_{id} + \boldsymbol{e}, s \cdot u + e_{2m+1} + \lfloor q/2 \rceil M)$.
- **Decrypting $(\boldsymbol{c}_1, c_2)$ addressed to user U_{id}.** Given ciphertext $(\boldsymbol{c}_1, c_2)$ and secret key T_{id}, compute $\boldsymbol{x} = T_{id}\boldsymbol{c}_1 \bmod q$ (which should be $T_{id}\boldsymbol{e}$ over R, as both T_{id} and $\boldsymbol{e}$ are small). Compute $\boldsymbol{e}' = T_{id}^{-1}\boldsymbol{x}$ over K. A candidate s' for s is obtained by dividing (in R_q) the first component of $\boldsymbol{c}_1 - \boldsymbol{e}'$ by the first component of $\boldsymbol{a}_0$. Now, Round each coefficient of $c_2 - s'u \bmod q$ to b if it is closer to $b\lfloor q/2\rceil$ than to $(1-b)\lfloor q/2 \rceil$. Return the obtained $M' \in R$ with coefficients in $\{0, 1\}$.

Fig. 1. A RLWE-based variant of the Agrawal et al IBE scheme.

The main differences between the IBE from Agrawal et al and the one from Figure 1 are as follows. First, in the master key generation, we use the structured variant of the Alwen-Ajtai-Peikert trapdoor generation algorithm [5,6], proposed in [42]. This provides a pair $(T, \boldsymbol{a}) \in R^{m\times m} \times R_q^m$ such that the distribution of $\boldsymbol{a}$ is within exponentially small statistical distance from uniform over R_q^m, the matrix T has full rank over K, it satisfies $T\boldsymbol{a} = \boldsymbol{0} \bmod q$, and the entries of T are polynomials with small coefficients. Using the ring variant of the recent algorithm from [30] may lead to a more efficient scheme in practice. Second, we encode identities as elements of the finite field $\mathbb{F}_q^{n/2}$. These are mapped to elements of R_q by using a diagonal encoding: We let $\mathcal{H}(id)$ denote the unique element of R_q that is congruent to id when reduced both modulo Φ_1 and Φ_2, i.e., its Chinese Remainder Theorem decomposition is $(id, id) \in (\mathbb{F}_{q^{n/2}})^2$. The property that allows for the security proof of [1] to carry over to our setting is that $\mathcal{H}(id) - \mathcal{H}(id') = \mathcal{H}(id - id')$ is invertible in R_q whenever $id \neq id'$. An alternative solution could be to let identities be the elements of R_q of degrees $< n/2$, as all such polynomials are necessarily non-zero modulo Φ_1 and Φ_2. Third, the plaintext space $\mathcal{P}$ is the set of polynomials in R with coefficients in $\{0, 1\}$, which allows for the encryption and decryption of n bits at once. Finally, the linear algebra operations performed in the decryption procedure can all be implemented in time quasi-linear in n, as these consist in small dimensional linear algebra tasks over the polynomial rings R_q and R (with coefficients of small magnitudes in the latter case).

Typical parameters for which correctness holds and for which the security proof of [1] carries over are $q = O(n^c)$, $m = O(\log q)$ and $\alpha = O(n^{-c'})$ for some constants c and c'.

Acknowledgments We thank Ron Steinfeld for very helpful discussions. Significant parts of the research corresponding to this work were undergone while the authors were visiting Macquarie University, whose hospitality is gratefully acknowledged. The authors were partly supported by the Australian Research Council Discovery Grant DP110100628.

References

1. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *Proc. of EUROCRYPT*, volume 6110 of *LNCS*, pages 553–572. Springer, 2010.
2. S. Agrawal, D. Boneh, and X. Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *Proc. of CRYPTO*, volume 6223 of *LNCS*, pages 98–115. Springer, 2010.
3. S. Agrawal, X. Boyen, V. Vaikuntanathan, P. Voulgaris, and H. Wee. Fuzzy identity based encryption from lattices, 2011. Available at `http://eprint.iacr.org/2011/414`.
4. S. Agrawal, D. M. Freeman, and V. Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. In *Proc. of ASIACRYPT*, volume 7073 of *LNCS*, pages 21–40. Springer, 2011.
5. M. Ajtai. Generating hard instances of the short basis problem. In *Proc. of ICALP*, volume 1644 of *LNCS*, pages 1–9. Springer, 1999.
6. J. Alwen and C. Peikert. Generating shorter bases for hard random lattices. *Theor. Comput. Science*, 48(3):535–553, 2011.
7. 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.
8. A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. To appear in the proceedings of EUROCRYPT'12. Available at `http://www.cc.gatech.edu/~cpeikert/pubs/prf-lattice.pdf`.
9. D. Boneh, R. Canetti, S. Halevi, and J. Katz. Chosen-ciphertext security from identity-based encryption. *SIAM J. Comput*, 36(5):1301–1328, 2007.
10. D. Boneh and B. Waters. Conjunctive, subset, and range queries on encrypted data. In *Proc. of TCC*, volume 4392 of *LNCS*, pages 535–554. Springer, 2007.
11. Z. Brakerski, C. Gentry, and V. Vaikuntanathan. Fully homomorphic encryption without bootstrapping, 2011. Available at `http://eprint.iacr.org/2011/277`.
12. Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *Proc. of FOCS*, pages 97–106. IEEE Computer Society Press, 2011.
13. Z. Brakerski and V. Vaikuntanathan. Fully homomorphic encryption from Ring-LWE and security for key dependent messages. In *Proc. of CRYPTO*, volume 6841 of *LNCS*, pages 505–524. Springer, 2011.
14. D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *Proc. of Eurocrypt*, volume 6110 of *LNCS*, pages 523–552. Springer, 2010.
15. J. Chen, H. W. Lim, S. Ling, H. Wang, and T. T. K. Nguyen. Revocable identity-based encryption from lattices, 2011. Available at `http://eprint.iacr.org/2011/583`.
16. J. von zur Gathen and J. Gerhardt. *Modern Computer Algebra, 2nd edition.* Cambridge University Press, 2003.
17. C. Gentry, S. Halevi, and N. Smart. Fully homomorphic encryption with polylog overhead. To appear in the proceedings of EUROCRYPT'12. Available at `http://eprint.iacr.org/2011/566`.
18. C. Gentry, S. Halevi, and V. Vaikuntanathan. A simple BGN-type cryptosystem from LWE. In *Proc. of EUROCRYPT*, volume 6110 of *LNCS*, pages 506–522. Springer, 2010.
19. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008.
20. 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.

21. S. Heyse, E. Kiltz, V. Lyubashevsky, C. Paar, and K. Pietrzak. An efficient authentication protocol based on Ring-LPN. To appear in the proceedings of FSE'12. Available at `http://www.uclouvain.be/crypto/ecrypt_lc11/static/post_proceedings.pdf`.
22. J. Katz, A. Sahai, and B. Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *Proc. of EUROCRYPT*, volume 4965 of *LNCS*, pages 146–162. Springer, 2008.
23. 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*, pages 7–26. Springer, 2011.
24. A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. Available at `http://perso.ens-lyon.fr/damien.stehle/MSIS.html/`.
25. S. Ling and D. Stehlé. A lattice-based traitor tracing scheme, 2012. Available at `http://perso.ens-lyon.fr/damien.stehle/LBTT.html`.
26. V. Lyubashevsky and D. Micciancio. Asymptotically efficient lattice-based digital signatures. In *Proc. of TCC*, volume 4948 of *LNCS*, pages 37–54. Springer, 2008.
27. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *Proc. of EUROCRYPT*, volume 6110 of *LNCS*, pages 1–23. Springer, 2010.
28. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings, 2011. Draft for the extended version of [27], dated 01/02/2011.
29. D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *Proc. of CRYPTO*, volume 6841 of *LNCS*, pages 465–484. Springer, 2011.
30. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. To appear in the proceedings of EUROCRYPT'12. Available at `http://www.cc.gatech.edu/~cpeikert/pubs/efftrap.pdf`.
31. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput*, 37(1):267–302, 2007.
32. C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. In *Proceedings of the 2007 IEEE Conference on Computational Complexity*, pages 333–346. IEEE Computer Society Press, 2007.
33. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.
34. C. Peikert. An efficient and parallel gaussian sampler for lattices. In *Proc. of CRYPTO*, volume 6223 of *LNCS*, pages 80–97. Springer, 2010.
35. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *Proc. of STOC*, pages 187–196. ACM, 2008.
36. K. Pietrzak. Subspace LWE, 2011. Available at `http://homepages.cwi.nl/~pietrzak/publications/SLWE.pdf`.
37. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proc. of STOC*, pages 84–93. ACM, 2005.
38. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
39. O. Regev. The learning with errors problem, 2010. Invited survey in CCC 2010, available at `http://www.cs.tau.ac.il/~odedr/`.
40. R. Sepahi, R. Steinfeld, and J. Pieprzyk. Lattice-based completely non-malleable PKE in the standard model (poster). In *Proc. of ACISP*, volume 6812 of *LNCS*, pages 407–411. Springer, 2011.
41. 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.
42. 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.
43. R. Steinfeld, S. Ling, J. Pieprzyk, C. Tartary, and H. Wang. NTRUCCA: How to strengthen NTRUEncrypt to chosen-ciphertext security in the standard model. Submitted for publication.

A Missing proofs

A.1 Proof of Lemma 5

This proof follows the same principle as the one of [38, Claim 3.9]. Using the Poisson summation formula, one obtains that the probability density function Y can be written as:

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

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

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

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

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

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

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

A.2 Proof of Lemma 8

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

If the sample $(\boldsymbol{a}, b)$ is uniform over $(R_q)^d \times \mathbb{T}_{R^*}$, then $(b + f \bmod R^*)$ is uniform in $\mathbb{T}_{R^*}$. Now, assume that $(\boldsymbol{a}, b)$ is distributed according to $A^{(q)}_{\boldsymbol{s}, \nu_{\boldsymbol{r}}}$: We have $b = \frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s} \rangle + e$, where $e \sim \nu_{\boldsymbol{r}}$. Since $\frac{1}{q}\langle \boldsymbol{a}, \boldsymbol{s} \rangle$ belongs to $\frac{1}{q}R^*$, we have $D_{(1/q)R^* - b, \boldsymbol{r}} = D_{(1/q)R^* - e, \boldsymbol{r}}$. Then, we use [34, Th. 3.1] (in the same fashion as in the proof of [20, Le. 2]), to state that sampling e from $\nu_{\boldsymbol{r}}$ and then setting $e' = e + f$ with f sampled from $D_{(1/q)R^* - e, \boldsymbol{r}}$ gives that the distribution of e' is statistically close to $D_{(1/q)R^*, \sqrt{2}\boldsymbol{r}}$. We conclude that in this case, the transformation returns a sample of $A^{(q)}_{\boldsymbol{s}, D_{(1/q)R^*, \sqrt{2}\boldsymbol{r}}}$. □

A.3 Proof of Lemma 9

We are given samples from a distribution D that is either the uniform distribution in $(R_q)^d \times \mathbb{T}_{R^*}$, or from $A^{(q)}_{\boldsymbol{s}, D_{(1/q)R^*, \boldsymbol{r}}}$.

In a first stage, we take several samples $(\boldsymbol{a}, b)$ from D and construct a set of d pairs $\{(\boldsymbol{a}_i, b_i)\}$ such that the $\boldsymbol{a}_i$'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 $\boldsymbol{a}_i$'s. This can be observed by considering the CRT components of $R_q \simeq (\mathbb{F}_{q^k})^{\frac{n}{k}}$ independently: An equivalent condition is that the n/k matrices corresponding to each component are invertible over the corresponding finite field. We define $\overline{\boldsymbol{A}} = (\boldsymbol{a}_1^T, \ldots, \boldsymbol{a}_d^T)$ and $\overline{\boldsymbol{b}} = (b_1, \ldots, b_d)^T$. By construction, the map $\boldsymbol{y} \mapsto \overline{\boldsymbol{A}}\boldsymbol{y}$ is a bijection of $(R_q)^d$, and if $D = A^{(q)}_{\boldsymbol{s}, D_{(1/q)R^*, \boldsymbol{r}}}$ then we have $\overline{\boldsymbol{b}} = \frac{1}{q}\left(\overline{\boldsymbol{A}}\boldsymbol{s} + \overline{\boldsymbol{x}} \right)$, where $\overline{\boldsymbol{x}}$ is sampled from $D_{(R^*)^d, q\boldsymbol{r}}$.

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

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

□

B A fast authentication scheme

The authentication scheme from Figure 2 is a direct adaptation to the RLWE setting of the authentication scheme from [21] based on the Ring version of the Learning Parity with Noise problem (LPN). The latter can itself be naturally interpreted as an efficient variant of the authentication scheme from [23], which can be instantiated to have its security rely on the presumed hardness of either LWE or LPN.

The function $\mathcal{H}$ is the same as in the identity-based encryption scheme from Section 4.

- **Key generation.** The shared key is a uniformly sampled pair $(s_1, s_2) \in (R_q)^2$.
- **Phase 1.** The Verifier samples v uniformly in $\mathbb{F}_{q^{n/2}}$ and sends $\mathcal{H}(v)$ to the prover.
- **Phase 2.** The Prover samples r uniformly among the invertible elements of R_q and e from the RLWE noise distribution $\phi \hookleftarrow \overline{\Upsilon}_{\alpha,q}$. The Prover then replies to query $\mathcal{H}(v)$ by sending (r, z) to the Verifier, with $z = r(s_1\mathcal{H}(v) + s_2) + e$.
- **Phase 3.** After receiving (r, z) from the Prover, the Verifier accepts if and only if r is invertible and $z - r(s_1\mathcal{H}(v) + s_2) \bmod q$ has sufficiently small coefficients.

Fig. 2. A RLWE-based variant of the Heyse et al authentication scheme.