JYX JYVÄSKYLÄN YLIOPISTO UNIVERSITY OF JYVÄSKYLÄ

This is a self-archived version of an original article. This version may differ from the original in pagination and typographic details.

Author(s): Grover, Charles; Mendelsohn, Andrew; Ling, Cong; Vehkalahti, Roope

Title: Non-commutative Ring Learning with Errors from Cyclic Algebras

Year: 2022

Version: Published version

Copyright: © The Author(s) 2022

Rights: CC BY 4.0

Rights url: https://creativecommons.org/licenses/by/4.0/

Please cite the original version:

Grover, C., Mendelsohn, A., Ling, C., & Vehkalahti, R. (2022). Non-commutative Ring Learning with Errors from Cyclic Algebras. Journal of Cryptology, 35(3), Article 22. https://doi.org/10.1007/s00145-022-09430-6

J Cryptol (2022) 35:22
https://doi.org/10.1007/s00145-022-09430-6

Journal of CRYPTOLOGY

Research Article

Non-commutative Ring Learning with Errors from Cyclic Algebras

Charles Grover · Andrew Mendelsohn · Cong Ling
Imperial College London, London, UK
c.grover15@imperial.ac.uk
andrew.mendelsohn18@imperial.ac.uk
cling@ieee.org

Roope Vehkalahti
Department of Mathematics and Statistics, University of Jyväskylä, 40014 Jyväskylä, Finland
roope.i.vehkalahti@jyu.fi

Communicated by Damien Stehlé

Received 19 November 2020 / Revised 18 May 2022 / Accepted 19 May 2022

Abstract. The Learning with Errors (LWE) problem is the fundamental backbone of modern lattice-based cryptography, allowing one to establish cryptography on the hardness of well-studied computational problems. However, schemes based on LWE are often impractical, so Ring LWE was introduced as a form of 'structured' LWE, trading off a hard to quantify loss of security for an increase in efficiency by working over a well-chosen ring. Another popular variant, Module LWE, generalizes this exchange by implementing a module structure over a ring. In this work, we introduce a novel variant of LWE over cyclic algebras (CLWE) to replicate the addition of the ring structure taking LWE to Ring LWE by adding cyclic structure to Module LWE. We show that the security reductions expected for an LWE problem hold, namely a reduction from certain structured lattice problems to the hardness of the decision variant of the CLWE problem (under the condition of constant rank d). As a contribution of theoretic interest, we view CLWE as the first variant of Ring LWE which supports non-commutative multiplication operations. This ring structure compares favorably with Module LWE, and naturally allows a larger message space for error correction coding.

Keywords. Algebraic number theory, Lattices, Learning with errors, Non-commutative algebra, Post-quantum cryptography.

1. Introduction

With the predicted advent of quantum computers compromising the bulk of existent cryptographic constructions, lattice-based cryptography has emerged as a promising foundation for long term security. In particular, the Learning with Errors (henceforth

© The Author(s) 2022
Published online: 13 July 2022

LWE) problem introduced in [42], as well as its variants over rings (RLWE) [27] and modules (MLWE) [22], provides a natural intermediate step to base cryptographic hardness on lattice short vector problems in a post-quantum setting. Indeed, second round submissions to the NIST post-quantum standardization process such as NewHope [3] and KYBER [5] rely on the hardness of LWE variants. Cryptography based on the classical LWE problem is typically somewhat impractical, in part due to large key sizes. To solve this, the ring variant was introduced as a way to provide extra structure in LWE to trade a potential loss of security for an increase in efficiency. MLWE generalizes ring and classical LWE, providing a smoother transition between security and efficiency than the binary option presented by ring or classical LWE. The flexibility of MLWE is highly desirable in practice, as demonstrated by third-round NIST finalists KYBER and SABER, both based on MLWE [1].

Conceptually, one may view all these problems as variations on a single problem. The (search) LWE problem tasks a solver with recovering a secret vector $\mathbf{s} \in \mathbb{Z}_q^n$ from a collection of pairs $(\mathbf{a}_i, b = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i)$, where $\langle \cdot, \cdot \rangle$ denotes the inner product, each $\mathbf{a}_i \in \mathbb{Z}_q^n$ is uniformly random and the e_i's are small random errors. In practice, we view this collection of equations in matrix–vector form:

$$A\mathbf{s} + \mathbf{e} = \mathbf{b},$$

where all operations and entries are over $\mathbb{Z}_q$ and the challenge is to recover $\mathbf{s}$ from $A, \mathbf{b}$. A popular ring variant replaces $A, \mathbf{s}, \mathbf{e}$ with elements a, s, e from the ring $R_q := \frac{\mathbb{Z}_q[x]}{x^n+1}$, requiring the solver to obtain s from samples $a_i \cdot s + e_i$. For power-of-two n this can be expressed in matrix–vector form by considering the matrix $\mathrm{rot}(a)$, the negacyclic matrix obtained from the coefficients of a. Explicitly, for $a = a_0 + a_1 x + \ldots + a_{n-1} x^{n-1}$ and bold faced letters denoting coefficient vectors, a sample from the RLWE distribution takes the form:

$$\begin{pmatrix} a_0 & -a_{n-1} & \ldots & -a_1 \\ a_1 & a_0 & \ldots & -a_2 \\ \vdots & \vdots & \ddots & \vdots \\ a_{n-1} & a_{n-2} & \ldots & a_0 \end{pmatrix} \mathbf{s} + \mathbf{e} = \mathbf{b}$$

where once again operations and entries are over $\mathbb{Z}_q$. This is exactly a structured version of the classical LWE problem, where the uniformly random matrix A has been replaced by the negacyclic matrix $\mathrm{rot}(a)$. Of course, this should be no harder to solve, yet no substantial progress has been made in using the structure of $\mathrm{rot}(a)$ to solve the problem efficiently. We can extend this matrix–vector view to MLWE as well. An MLWE instance takes place in a module M of dimension d over R_q, such that a solver has to recover $\mathbf{s} \in M$ from a collection of pairs $(\mathbf{a}_i, \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i)$ where $\mathbf{a}_i$ is a uniformly random element of M and each e_i is a small random element of R_q. A collection of such pairs can be viewed as $A\mathbf{s} + \mathbf{e} = \mathbf{b}$, where the ambient space $\mathbb{Z}_q$ has been replaced by R_q, e.g., with d samples:

$$\begin{pmatrix} a_{1,1} & a_{1,2} & \dots & a_{1,d} \\ a_{2,1} & a_{2,2} & \dots & a_{2,d} \\ \vdots & \vdots & \ddots & \vdots \\ a_{d,1} & a_{d,2} & \dots & a_{d,d} \end{pmatrix} \mathbf{s} + \mathbf{e} = \mathbf{b}$$

where all operations are over R_q and each $a_{i,j}$ is uniformly random. Of course, we could extend this to have operations over $\mathbb{Z}_q$ by applying the rot$(\cdot)$ operation coordinatewise, to obtain a structured LWE instance in dimension nd.

An advantage of these structured matrices is that they allow for streamlined storage and operations. For example, storing a uniformly random matrix A requires one to store all n^2 of its entries, but rot(a) requires a factor n less memory since one need only store its first column. Equivalently, one RLWE sample generates n LWE samples while reducing the storage space and key sizes. Multiplication can also be speeded up by using the Chinese Remaindering Theorem (CRT) or other techniques.

This concept of improving efficiency by adding structure motivates this work; can we perform an analog of the transformation taking an LWE matrix A to an RLWE matrix rot(a) for the module M? We solve this by constructing a new variant of the LWE problem over a certain non-commutative space known as a *cyclic algebra*. In recent years, cyclic algebras have received significant attention in the field of coding theory (see, e.g., [25,32,44]) due to the particular nature of the matrix lattices they induce, and we view them as a suitable option for defining an LWE problem over a non-commutative ring. Though some efforts have been made to construct non-commutative LWE problems, for example [8,16], the majority of non-commutative cryptography has relied on group theoretic constructions, whose underlying hard problems are often less robust than those of lattice cryptography. Somewhat informally, for a cyclic algebra $\mathcal{A}$ and well-chosen parameters there exists an automorphism θ of R_q and a $\gamma \in R_q$ such that an LWE style sample $a \cdot s + e$ over $\mathcal{A}$ can be written in matrix–vector form

$$\begin{pmatrix} a_0 & \gamma\theta(a_{d-1}) & \gamma\theta^2(a_{d-2}) & \dots & \gamma\theta^{d-1}(a_1) \\ a_1 & \theta(a_0) & \gamma\theta^2(a_{d-1}) & \dots & \gamma\theta^{d-1}(a_2) \\ a_2 & \theta(a_1) & \theta^2(a_0) & \dots & \gamma\theta^{d-1}(a_3) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{d-1} & \theta(a_{d-2}) & \theta^2(a_{d-3}) & \dots & \theta^{d-1}(a_0) \end{pmatrix} \mathbf{s} + \mathbf{e} = \mathbf{b}$$

where all entries and operations are now over R_q. Though more complex than the transformation taking LWE to RLWE this fulfills our goal of providing a structured version of MLWE, since we have replaced the uniformly random matrix A over R_q with a structured matrix which we denote $\phi(a)$ that requires a factor of d less storage. Of course, by applying the rot$(\cdot)$ operation coordinatewise, one can extend this to a high-dimensional version of the LWE problem, now with two sets of structure lying on top of each other.

1.1. *Contributions and Methodology*

The main novel contribution of this work is a definition of Cyclic Algebra LWE (CLWE), together with justifications for its construction and a polynomial time reduction from short vector problems over matrix lattices induced by two-sided ideals in the maximal order of a cyclic algebra to CLWE, establishing its security on the assumption that such problems are hard. As in [27], the algorithm bases the security of CLWE on short vector problems over two-sided ideal lattices in $\mathcal{A}$; similarly to ideal lattices in K, these have some extra underlying structure that might make computational problems easier. However, we leave the relative complexity of these problems an open area of investigation.

CLWE represents a middle ground between RLWE and MLWE. Cyclic algebras are equipped with a proper ring multiplication which preserves the dimension of the lattice. Specifically, we consider the following advantages of our CLWE construction:

- Efficiency. CLWE can be seen a structured variant of MLWE. Assuming for simplicity that the public key in LWE-based schemes is a sample $(A, \mathbf{b})$, a public key generated as $A = \mathrm{rot}(\phi(a))$ requires only as much storage as that of an equivalent dimension RLWE public key.[1] On the negative side, one should note that we do not know currently how to construct CLWE instances of arbitrary dimension, which might have an impact on concrete efficiency of the schemes.
- Security. Recent works on quantum attacks on related ideal lattice problems (e.g., [10,14,17,18] amongst others) require that the underlying group, in this case the unit group of $\mathcal{O}_K$, is commutative, see, e.g., [20], which is untrue for a non-commutative algebra. We conjecture that the security level is higher than RLWE, but welcome further cryptanalysis. We actively avoid known attacks on previous attempts to create structured MLWE (see Sect. 3.2). We remark that solving ideal-SVP in a number field is not known to impact the security of RLWE. Moreover, there are currently no known algorithms solving RLWE faster than MLWE for similar parameter sets (either theoretically or practically). It is even known from [2] that for some specific choices of parameters, RLWE is asymptotically no easier than MLWE.
- Decryption failure rates. The scalar multiplication of MLWE is dimension-lossy. In other words, the message space of MLWE is restricted in R_q, whose dimension is smaller than that of the module lattice. It leaves less room for error correction coding in MLWE-based schemes (e.g., a KYBER instance for a key size of 256 within R_q of dimension 256). In contrast, the dimension of the message space of CLWE is that of the (non-commutative) ring, which is higher by a factor of d. Thus, it accommodates better error correction coding (see Sect. 5.2), and low decryption failure rates are desired under chosen ciphertext attacks (CCA). Even trivial repetition coding can dramatically reduce decryption failure rates (e.g., NewHope)[2]

[1] In practice, a seed is often used to generate the matrix A, which, however, requires a pseudorandom generator under the random oracle model. By contrast, CLWE does not require the random oracle model. Moreover, certain applications do not permit the use of a seed, *e.g.*, pseudorandom functions [7].

[2] The same result could be obtained in MLWE by increasing the public key and ciphertext sizes by a factor 2: instead of considering an MLWE sample $(A, \mathbf{b})$ with a vector $\mathbf{b}$, one could consider a square matrix B, whose columns correspond to independent LWE samples using the same matrix A.

Our search-to-decision reduction only holds for one choice of modulus q (once the algebra has been fixed) and structured modules of constant rank d. This issue needs to be remedied in the future.

1.2. *Related Work and Organization*

This work is related to a number of different areas: lattice-based cryptography, information theory and number theory.

In lattice-based cryptography, an alternative construction for structured module LWE, called multivariate-RLWE, was presented in [35,36], where they tensor product two (or more) number fields in order to provide a structured module matrix. However, an efficient implementation of [35] was attacked in [12], together with a warning about taking care when putting structure on a module. In short, [12] attacks certain instances of multivariate-RLWE by providing a homomorphism to some underlying subfield K, dramatically reducing the dimension of the lattice problem to be attacked. Fortunately for this work, a somewhat technical condition on the choice of γ known as the *non-norm condition* precludes such a homomorphism existing to reduce the dimension of CLWE (see Sect. 3.2). It is worth pointing out that that their problem has been addressed in [36], and in fact this fix looks somewhat like our non-norm condition (*e.g.*, unlike the original version, full rank is maintained in [36]).

This paper is inspired by the abundant literature of space-time coding based on cyclic division algebras (see the monographs [9,32] and references therein). On a high level, our construction is reminiscent of multiblock space-time codes [21,23], with the caveat of scaling up the number of blocks to make the codes practically undecodable. In the context of space-time coding, our construction generalizes [21] and offers greater flexibility in the code parameters (the number of blocks vs. the number of antennas). Multiblock space-time codes have been used in [25] to achieve information-theoretic security over wiretap channels, as opposed to computational security in a classic cryptographic setting of this paper. There is a major difference between the roles of cyclic algebras in coding and cryptography, though: the primary concern for coding is the non-vanishing determinant (NVD), while the non-commutative ring structure becomes crucial for cryptography. For efficient multiplication of elements in a cyclic algebra, we heavily rely on the CRT technique of [33].

We present two approaches (subfields and compositum fields) to the construction of novel cyclic division algebras, which enlarge the pool of algebras and may find other applications. Specifically, our proof that the natural order of the family of cyclic division algebras constructed in Theorem 2 (including those in [21]) is in fact maximal, is an original contribution.

The rest of this paper is organized as follows. In Sect. 2 we provide necessary background material on lattices, number fields, and cyclic algebras. In Sect. 3 we provide a definition and discussion of CLWE, together with novel constructions of cyclic division algebras for the CLWE problem. In Sect. 4 we provide a reduction from structured lattice problems to search CLWE, as well as a search-worst case decision reduction for CLWE. In Sect. 5 we show a sample CLWE cryptosystem and provide an estimate of its asymptotic operation complexity. Finally, the paper is concluded in Sect. 6 with a

discussion of open problems. For a smooth flow of the main text, certain proofs, sideline discussions and technical details are deferred to appendices.

2. Preliminaries

2.1. *Lattices*

A lattice is a discrete additive subgroup of a vector space V. If V has dimension n a lattice $\mathcal{L}$ can be viewed as the set of all integer linear combinations of a set of linearly independent vectors $B = \{\mathbf{b}_1, \ldots, \mathbf{b}_k\}$ for some $k \leq n$, written $\mathcal{L} = \mathcal{L}(B) = \{\sum_{i=1}^{k} z_i \mathbf{b}_i : z_i \in \mathbb{Z}\}$. If $k = n$ we call the lattice full-rank, and we will only consider lattices of full-rank. We can extend this notion of lattices to matrix spaces by stacking the columns of a matrix. We recall two standard lattice definitions.

Definition 1. Given a lattice $\mathcal{L}$ in a space V endowed with a metric $\|\cdot\|$, the minimum distance of $\mathcal{L}$ is defined as $\lambda_1(\mathcal{L}) = \min_{\mathbf{v} \in \Lambda/\{0\}} \|\mathbf{v}\|$. Similarly, $\lambda_n(\mathcal{L})$ is the minimum length of a set of n linearly independent vectors, where the length of a set of vectors $\{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ is defined as $\max_i(\|\mathbf{x}_i\|)$.

Definition 2. Given a lattice $\mathcal{L} \subset V$, where V is endowed with an inner product $\langle \cdot, \cdot \rangle$, the dual lattice $\mathcal{L}^*$ is defined $\mathcal{L}^* = \{\mathbf{v} \in V : \langle \mathcal{L}, \mathbf{v} \rangle \subset \mathbb{Z}\}$.

2.2. *Gaussian Distributions*

Definition 3. For a vector space V with norm $\|\cdot\|$ and an $r > 0$, we define the Gaussian function $\rho_r : V \to (0, 1]$ by $\rho_r(\mathbf{x}) = \exp(-\pi \|\mathbf{x}\|/r^2)$.

We can use this function to define the spherical Gaussian distribution D_r over V, which outputs $\mathbf{v}$ with probability proportional to $\rho_r(\mathbf{v})$. Similarly, we can sample an elliptical Gaussian $D_{\mathbf{r}}$ in a basis $\mathbf{b}_1, \ldots, \mathbf{b}_n$ of V, for $\mathbf{r} = (r_1, \ldots, r_n)$ a vector of positive reals, by sampling $x_1, \ldots, x_n$ independently from the one-dimensional Gaussian distributions D_{r_i} and outputting $\sum_{i=1}^{n} x_i \mathbf{b}_i$.

When sampling a Gaussian over a lattice $\mathcal{L}$, we will use the discrete form of the Gaussian distribution. We define the distribution $D_{\Lambda,r}$ over Λ by outputting $\mathbf{x}$ with probability $\frac{\rho_r(\mathbf{x})}{\rho_r(\mathcal{L})}$ for each $\mathbf{x} \in \mathcal{L}$. This version of the discrete Gaussian is centered at 0, which in general need not be the case.

An important lattice quantity, known as the smoothing parameter, was introduced in [31]. The motivation for the name is provided by Sect. 1 following the definition.

Definition 4. For a lattice $\mathcal{L}$ and $\varepsilon > 0$, the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ is defined as the smallest $r > 0$ satisfying $\rho_{1/r}(\mathcal{L}^*/\{\mathbf{0}\}) \leq \varepsilon$.

The following is a special case of [31], Lemma 4.1.

Lemma 1. *For a lattice $\mathcal{L}$ over $\mathbb{R}^n$, $\varepsilon > 0, r \geq \eta_\varepsilon(\mathcal{L})$, and $\boldsymbol{x} \in \mathbb{R}^n$, the statistical distance between $(D_r + \boldsymbol{x}) \mod \mathcal{L}$ and the uniform distribution modulo $\mathcal{L}$ is bounded above by $\varepsilon/2$. Equivalently, $\rho_r(\mathcal{L} + \boldsymbol{x}) \in [\frac{1-\varepsilon}{1+\varepsilon}, 1] \cdot \rho_r(\mathcal{L})$.*

We introduce well-known lemmas used to relate the smoothing parameter to standard lattice properties. The first comes from [6], the second from [40].

Lemma 2. *For a lattice $\mathcal{L}$ of dimension n and $c \geq 1$, it holds that $c\sqrt{n}/\lambda_1(\mathcal{L}^*) \geq \eta_\varepsilon(\mathcal{L})$ for $\varepsilon = \exp(-c^2 n)$.*

Lemma 3. *For a lattice $\mathcal{L}$ and $\varepsilon \in (0, 1)$, it holds that $\eta_\varepsilon(\mathcal{L}) \leq \frac{\sqrt{\log(1/\varepsilon)/\pi}}{\lambda_1(\mathcal{L}^*)}$.*

2.3. *Algebraic Number Theory*

Definition 5. A number field K is a finite degree extension of the rationals $\mathbb{Q}$. Typically, we define a number field by adjoining some algebraic element $\alpha \in \mathbb{C}$ and set $K = \mathbb{Q}(\alpha)$. The degree of K refers to its degree as a field extension.

To define a cyclic algebra, we will need to take an additional extension of K. In particular, we will need the extension to be Galois over K, defined as follows.

Definition 6. Let L/K be an extension of number fields of dimension d. The Galois group of L over K is the group $\mathrm{Aut}(L/K)$ of automorphisms of L that fix K. We say that the extension is Galois if the subfield of L fixed by $\mathrm{Aut}(L/K)$ is exactly K.

We define a cyclic Galois extension L/K to be a Galois extension such that the Galois group of L over K is the cyclic group generated by some element θ of degree $d := [L : K]$. Finally, we require the ring of integers of a number field.

Definition 7. Given a number field K, its ring of integers $\mathcal{O}_K$ is the ring consisting of those elements of K whose minimal polynomial over $\mathbb{Q}$ lie in $\mathbb{Z}[x]$.

It is easy to check that if L/K is an extension of number fields then $\mathcal{O}_L \cap K = \mathcal{O}_K$.

2.3.1. *The Canonical Embedding*

Let $K = \mathbb{Q}(\alpha)$ be a number field of degree n. It is a well-known fact that there are exactly n distinct ring embeddings $\sigma_i : K \to \mathbb{C}$. These embeddings correspond to the n distinct injective ring homomorphisms mapping α to the roots of its minimum polynomial f. We split these embeddings and say that there are r_1 real embeddings (whose image lie in $\mathbb{R}$) and r_2 conjugate pairs of complex embeddings (the complex embeddings come in pairs since complex roots of f occur in conjugate pairs), such that $r_1 + 2r_2 = n$. The standard convention is to order the embeddings such that the r_1 real embeddings come first and the complex embeddings are arranged such that $\sigma_{r_1+j} = \overline{\sigma_{r_1+r_2+j}}$ for $1 \leq j \leq r_2$.

Definition 8. Let $K = \mathbb{Q}(\alpha)$ be a number field of degree $n = r_1 + 2r_2$. The canonical embedding σ is the ring homomorphism $\sigma : K \to \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ defined by

$$\sigma(x) = (\sigma_1(x), \ldots, \sigma_n(x)).$$

Formally, σ maps into the space

$$H = \{(x_1, \ldots, x_n) \in \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2} \mid x_{r_1+r_2+j} = \overline{x_{r_1+j}} \ \forall 1 \leq j \leq r_2\} \subset \mathbb{C}^n,$$

which is isomorphic to $\mathbb{R}^n$ as an inner product space.

We can equip H with the orthonormal basis $\{\mathbf{h}_i\}$, where $\mathbf{h}_i = \mathbf{e}_i$ for $1 \leq i \leq r_1$ and $\mathbf{h}_j = \frac{1}{\sqrt{2}}(\mathbf{e}_j + \mathbf{e}_{j+r_2}), \mathbf{h}_{j+r_2} = \frac{\sqrt{-1}}{\sqrt{2}}(\mathbf{e}_j - \mathbf{e}_{j+r_2})$ for $r_1 < j \leq r_1 + r_2$, and use the well-defined ℓ_p norm induced by viewing H as a subset of $\mathbb{C}^n$. Observe that multiplication in K maps to coordinatewise multiplication in H. The ℓ_2 norm on H allows us to efficiently sample a Gaussian distribution $D_{\mathbf{r}}$ over K by sampling such a Gaussian coordinatewise over H, although technically this distribution is over the field tensor product $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R} \cong H$. Furthermore, it satisfies the property that for any $x \in K_{\mathbb{R}}$ we have the equality of distributions $x \cdot D_{\mathbf{r}}$ and $D_{\mathbf{r}'}$, where $r_i' = r_i \cdot |\sigma_i(x)|$. When we have an extension of number fields L/K, we will denote their respective canonical embeddings σ_L and σ_K as maps into H_L and H_K to avoid confusion.

2.3.2. *Relative Embeddings*

In the case of an extension L of a number field K it is sometimes more convenient to apply a different order on its embeddings induced by extending embeddings of K to those of L. Given a tower $L/K/\mathbb{Q}$ where K has degree n and L has degree d over K, there are precisely n embeddings $\sigma_1, \ldots, \sigma_n$ of K into $\mathbb{C}$. Assuming $L/\mathbb{Q}$ is Galois, each of these can be extended to an embedding $\alpha_i : L \to \mathbb{C}$ such that $\alpha_i|_K = \sigma_i$. However, these extensions are not unique, and it is easy to see that there are $[L : K] = d$ choices for each α_i. In particular, in the case where L/K is a cyclic extension with Galois group generated by θ it holds that the composite automorphisms $\alpha_i \circ \theta^j(\cdot)$, $1 \leq j \leq d$, run through the d choices of α_i. Hence, for a fixed choice of $\alpha_1, \ldots, \alpha_n$ the nd automorphisms of L can each be uniquely represented by some $\alpha_i \circ \theta^j(\cdot)$, which we denote by $\alpha_i^j(\cdot)$, $1 \leq i \leq n$, $1 \leq j \leq d$. Given the usual ordering of embeddings of K, this induces two systematic orderings on the embeddings of L by running through either the i or j coordinates first.

2.4. *Cyclic Algebras*

Definition 9. Let K be a number field with degree n, and let L be a Galois extension of K of degree d such that the Galois group of L over K is cyclic of degree d, $\mathrm{Gal}(L/K) = \langle\theta\rangle$. For nonzero $\gamma \in K$ we define the resulting cyclic algebra

$$\mathcal{A} = (L/K, \theta, \gamma) := L \oplus uL \oplus \ldots \oplus u^{d-1}L$$

$$K \xrightarrow[\text{Degree } d]{\text{Cyclic Extension}} L \xrightarrow[\text{Degree } d]{\text{Attach } u} \mathcal{A}$$

Fig. 1. Structure of a cyclic algebra.

where $\oplus$ denotes the direct sum, $u \in \mathcal{A}$ is some auxiliary generating element of $\mathcal{A}$ satisfying the additional relations $xu = u\theta(x), \forall x \in L$ and $u^d = \gamma$. We will call d the degree of the algebra $\mathcal{A}$. We call such an algebra a division algebra if every element $a \in \mathcal{A}$ has an inverse $a^{-1} \in \mathcal{A}$ such that $aa^{-1} = 1$.

The relations among K, L and $\mathcal{A}$ are illustrated in Fig. 1. In fact, every central simple algebra over a number field is cyclic.

Since θ fixes K, the center of the cyclic algebra is precisely K. Oftentimes the condition $\gamma \in K$ is replaced by the stronger condition $\gamma \in \mathcal{O}_K$, and we will use this condition in our work to guarantee the existence of a certain subring known as the natural order. Note that the division property does not hold for arbitrary γ, and such algebras are not always easy to construct, which we will discuss later in this section.

We present a matrix representation of elements of $\mathcal{A}$ which proves useful for computing multiplication in cyclic algebras. We can naturally view an element $a \in \mathcal{A}$ as an d-dimensional vector $\mathrm{Vec}(a)$ over L, in which case we can view left multiplication of elements as matrix–vector operations. This is done by defining the map $\phi : \mathcal{A} \to M_{d\times d}(L)$, where for $x = x_0 + ux_1 + ... + u^{d-1}x_{d-1} \in \mathcal{A}$ with each $x_i \in L$,

$$\phi(x) = \begin{pmatrix} x_0 & \gamma\theta(x_{d-1}) & \gamma\theta^2(x_{d-2}) & \dots & \gamma\theta^{d-1}(x_1) \\ x_1 & \theta(x_0) & \gamma\theta^2(x_{d-1}) & \dots & \gamma\theta^{d-1}(x_2) \\ x_2 & \theta(x_1) & \theta^2(x_0) & \dots & \gamma\theta^{d-1}(x_3) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x_{d-1} & \theta(x_{d-2}) & \theta^2(x_{d-3}) & \dots & \theta^{d-1}(x_0) \end{pmatrix}.$$

We call this mapping a left regular representation of $\mathcal{A}$, because it holds for any $a, b \in \mathcal{A}$ that $\phi(a)\mathrm{Vec}(b) = \mathrm{Vec}(ab)$, and that $\phi(ab) = \phi(a) \cdot \phi(b)$. In the case where $\mathcal{A}$ is a division algebra it follows that each $\phi(a)$ is an invertible matrix. Since θ is well defined on $L_{\mathbb{R}}$, we abuse notation and extend this map to $\phi : \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}} \to M_{d\times d}(L_{\mathbb{R}})$. We derive lattices from subrings of a cyclic algebra by vectorizing their images under ϕ.

Definition 10. Let $\mathcal{A} = (L/K, \theta, \gamma)$ be a cyclic division algebra. A $\mathbb{Z}$-order Λ in $\mathcal{A}$ is a finitely generated $\mathbb{Z}$-module such that $\Lambda \cdot \mathbb{Q} = \mathcal{A}$ and that Λ is a subring of $\mathcal{A}$ with the same identity element as $\mathcal{A}$. We call Λ maximal if there is no $\mathbb{Z}$-order Γ such that $\Lambda \subsetneq \Gamma \subsetneq \mathcal{A}$. Here, $\Lambda \cdot \mathbb{Q} = \{\sum_{i=1}^{m} a_i q_i : a_i \in \Lambda, q_i \in \mathbb{Q}, m \in \mathbb{Z}_{\geq 1}\}$.

Since we are only concerned with $\mathbb{Z}$-orders in this paper, we will just refer to them as orders.

Example 1. The ring of integers $\mathcal{O}_K$ of a number field K is the unique maximal order of a number field. In the case of cyclic algebras a maximal order is not necessarily unique.

An order of particular interest that we will use in our LWE construction is known as the *natural order*, defined as $\Lambda := \bigoplus_{i=0}^{d-1} u^i \mathcal{O}_L$. Unlike in the case of $\mathcal{O}_K$, this order is not necessarily maximal. (However, we are going to work with natural orders that are also maximal.) Note that in order for Λ to be closed under multiplication the element γ must lie in $\mathcal{O}_K$.

2.4.1. *Non-Norm Condition*

It is not a priori obvious whether well-defined cyclic division algebras or orders actually exist. As observed earlier, the existence of γ enforcing the division algebra condition is a key component in constructing such objects. Fortunately, it is sufficient for γ to satisfy the so-called non-norm condition [44].

Proposition 1. *The cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ of degree d is a division algebra if and only if none of the elements γ^t, $1 \leq t \leq d-1$, appears in $N_{L/K}(L)$, where $N_{L/K}$ represents the relative norm of L into K.*

In other words, this condition states that the lowest power of γ that is norm of some element of L, is γ^d.

2.4.2. *Order Ideals*

Analogous to the use of $\mathcal{O}_K$ ideals in RLWE, we will be interested in ideals of an order Λ of a cyclic division algebra $\mathcal{A}$. Although Λ is a ring, it is non-commutative—thus there are three types of ideals. A left (respectively right) ideal $\mathcal{I}$ of Λ is an additive subgroup of Λ such that for any $i \in \mathcal{I}$, $r \in \Lambda$, we have $r \cdot i \in \mathcal{I}$ (respectively $i \cdot r \in \mathcal{I}$). A two-sided ideal of Λ is an additive subgroup that is closed under left and right scaling by Λ, i.e., a right ideal that is also a left ideal. The sum and product of two ideals $\mathcal{I}, \mathcal{J}$ are defined as usual; $\mathcal{I}+\mathcal{J} = \{i+j : i \in \mathcal{I}, j \in \mathcal{J}\}$ and $\mathcal{I} \cdot \mathcal{J} = \{\sum_{l=1}^{m} i_l \cdot j_l : i_l \in \mathcal{I}, j_l \in \mathcal{J}, m \in \mathbb{N}\}$. In the case of two-sided ideals we have the standard notion of a fractional ideal; $\mathcal{I}$ is a fractional ideal of Λ if $c\mathcal{I} = \mathcal{J}$ for a two-sided ideal $\mathcal{J}$ and some $c \in K$. In the rest of this paper, a (fractional or integral) ideal is always restricted to be two-sided, unless otherwise stated.

We remark that the structure of the collection of two-sided ideals of the natural order is not as simple as those of $\mathcal{O}_K$, or indeed those of an arbitrary maximal order. In a maximal order, the group of two-sided ideals is a free abelian group generated by the prime (e.g., maximal) ideals [43, Theorem 22.10], from which one can deduce obvious definitions of inverse and coprime ideals. For a general order Λ, we define its prime ideals as its maximal two-sided ideals and the inverse of an ideal $\mathcal{I} \subset \Lambda$ is

$$\mathcal{I}^{-1} = \{x \in \mathcal{A} : \mathcal{I} \cdot x \cdot \mathcal{I} \subset \mathcal{I}\},$$

which lines up with the expected definition in the two-sided case (e.g., $\mathcal{I} \cdot \mathcal{I}^{-1} = \mathcal{I}^{-1} \cdot \mathcal{I} = \Lambda$).

For the case of the natural order, we do not have such a well-behaved ideal group, but a nice exposition is given in [33, Sect. 3]. In particular, for a two-sided ideal $\mathcal{I} \subset \Lambda$, $\mathcal{I} \cap \mathcal{O}_K$ is an ideal of $\mathcal{O}_K$. For an ideal $\mathcal{I} \subset \mathcal{O}_K$, $(\mathcal{I} \cdot \Lambda) \cap \mathcal{O}_K = \mathcal{I}$, from which it

follows that this intersection map is a surjection onto the ideals of $\mathcal{O}_K$. However, it is not in general an injection since several ideals of Λ may have the same intersection with $\mathcal{O}_K$. Since the ideals of Λ do not in general form a finitely generated abelian group, we define two ideals $\mathcal{I}, \mathcal{J}$ of Λ to be coprime if $\mathcal{I} + \mathcal{J} = \Lambda$.

Nonetheless, since the orders to be constructed in Theorem 2 are both natural and maximal, it will always hold for a two-sided ideal $\mathcal{I}$ that $\mathcal{I} \cdot \mathcal{I}^{-1} = \mathcal{I}^{-1} \cdot \mathcal{I} = \Lambda$ and $(\mathcal{I}^{-1})^{-1} = \mathcal{I}$. These properties will be required in the proofs of Lemmas 6 and 7.

2.4.3. *Some Useful Ideals*

For an order Λ we define the codifferent ideal

$$\Lambda^{\vee} = \{x \in \mathcal{A} : \mathrm{Tr}(x\Lambda) \subset \mathbb{Z}\}$$

where Tr refers to the reduced trace, defined $\mathrm{Tr}(a) := \mathrm{Tr}_{K/\mathbb{Q}}(\mathrm{Trace}(\phi(a)))$. Similarly, for an ideal $\mathcal{I}$ we define the dual ideal

$$\mathcal{I}^{\vee} = \{x \in \mathcal{A} : \mathrm{Tr}(x\mathcal{I}) \subset \mathbb{Z}\}.$$

Since the matrix trace satisfies $\mathrm{Trace}(AB) = \mathrm{Trace}(BA)$, this definition is two-sided. Note that the codifferent ideal and a general dual ideal may be fractional ideals rather than full ideals, and they satisfy the equality $\mathcal{I}^{\vee} = \Lambda^{\vee} \cdot \mathcal{I}^{-1}$ for any ideal $\mathcal{I}$.

We will also be interested in principal ideals, but must take more care with these than in commutative settings. For a central element $t \in K$, we can define simply $\langle t \rangle = t \cdot \Lambda$, the set of elements of Λ divisible by t. However, for a general t that does not lie in the center of Λ we need the slightly more complex definition

$$\langle t \rangle = \left\{ \sum_{i=1}^{m} r_i t s_i : r_i, s_i \in \Lambda, m \in \mathbb{N} \right\},$$

which can easily be seen to be a two-sided ideal, moreover the smallest one that contains t.

2.4.4. *Orders and Ideals as Lattices*

Any order Λ of a cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ has dimension nd^2 over $\mathbb{Z}$ and thus generates a lattice of dimension nd^2 over $\mathbb{Z}$. We will consider the following representation of these lattices, which extends naturally to ideals of orders as well. Consider an element $x = \bigoplus_{i=0}^{d-1} u^i x_i \in \Lambda$. We can consider x as a vector over H_L of dimension d by $\sigma_{\mathcal{A}}(x) := \{\sigma_L(x_0), \sigma_L(x_1), \ldots, \sigma_L(x_{d-1})\}$. Then, the collection $\sigma_{\mathcal{A}}(\Lambda)$ forms a lattice of dimension nd^2 over $\mathbb{Z}$. We will refer to this representation as the "module representation" and will sometimes double index the element x, denoting by $x_{i,j}$ the embedding $\sigma_j(x_i)$, and extend this notation in the obvious manner to the space $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$. Though this representation is conceptually simple, we remark that it has some drawbacks in the case where $|\sigma_i(\gamma)| \neq 1$ for some i when considering sizes of lattice elements; we will choose γ carefully in our constructions to remove this issue.

As in (R)LWE, we will need to sample Gaussian distributions over our ambient space in certain norms. In the case of RLWE, the continuous Gaussians are sampled in $K_{\mathbb{R}} \cong H$. Since a cyclic algebra $\mathcal{A}$ can be viewed as a d-dimensional algebra over L, we use the visualization from the previous subsection and sample our error distributions over $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$, which has the same structure as a vector space as ${H_L}^d$. For simplicity we restrict ourselves to the case when $|\sigma_i(\gamma)| = 1$ for each i. Although this is a strong condition on γ it holds in the case where it is a root of unity, which we will enforce later.

We consider the norm of an element of $\mathcal{A}$ to be equal to the norm of the corresponding module element in L^d of dimension nd^2 used in [22], e.g., $\|x\| = \|(\sigma_L(x_0), \sigma_L(x_1), \ldots, \sigma_L(x_{d-1}))\|_2$ for $x = x_0 + ux_1 + \ldots + u^{d-1}x_{d-1} \in \mathcal{A}$. It is straightforward to check that this is indeed a norm in the case where $|\sigma_i(\gamma)| = 1$ for each i, since γ is fixed under θ and multiplying by γ does not change the norm of an entry of σ_L. In fact, if $|\sigma_i(\gamma)| = 1$ for each i, $\|x\|$ is equivalent to a representation in Frobenius norm $\|\cdot\|_F$ of matrices:

$$\|x\|^2 = \sum_{\sigma \in \sigma_K} \|\sigma(\phi(x))\|_F^2$$

where σ, when applied to $\phi(x)$, is a short notation of its extension to L. We have

$$\begin{aligned}
\|xy\|^2 &= \sum_{\sigma \in \sigma_K} \|\sigma(\phi(xy))\|_F^2 = \sum_{\sigma \in \sigma_K} \|\sigma(\phi(x))\sigma(\phi(y))\|_F^2 \\
&\leq \sum_{\sigma \in \sigma_K} \|\sigma(\phi(x))\|_F^2 \|\sigma(\phi(y))\|_F^2 \\
&\leq \sum_{\sigma \in \sigma_K} \|\sigma(\phi(x))\|_F^2 \sum_{\sigma \in \sigma_K} \|\sigma(\phi(y))\|_F^2 = \|x\|^2 \|y\|^2
\end{aligned}$$

where the first inequality is due to the sub-multiplicativity of Frobenius norm. In this case, the norm is sub-multiplicative.

It is clear that this norm extends to any $y \in \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ in a natural manner. Now that we have defined a norm, it is easy to define a Gaussian distribution $D_{\mathbf{r}}$ on $\mathcal{A}$, or its discrete analogue on Λ by sampling over the module ${L_{\mathbb{R}}}^d$.

2.4.5. *The CRT*

In this subsection we state the CRT for order ideals, and deduce some important consequences. We note that the following lemmas are merely adaptations of those in [27, Sect. 2.3.8] extended to the case of cyclic algebras. The first is just the CRT.

Lemma 4. *Let $\mathcal{I}_1, \ldots, \mathcal{I}_r$ be pairwise coprime ideals of an order Λ of a cyclic algebra $\mathcal{A}$, and let $\mathcal{I} = \prod_{i=1}^r \mathcal{I}_i$. Then, the natural map $\Lambda \to \bigoplus_{i=1}^r (\Lambda/\mathcal{I}_i)$ induces an isomorphism $\Lambda/\mathcal{I} \to \bigoplus_{i=1}^r (\Lambda/\mathcal{I}_i)$.*

We call a CRT basis for a set of coprime order ideals $\mathcal{I}_1, \ldots, \mathcal{I}_r$ a basis $C = \{c_1, \ldots, c_r\}$ of elements of Λ satisfying $c_i = 1 \mod \mathcal{I}_i, c_i = 0 \mod \mathcal{I}_j$ for $i \neq j$.

Lemma 5. *Given pairwise coprime ideals $\mathcal{I}_1, \ldots, \mathcal{I}_r$ of an order Λ, there is a deterministic polynomial time algorithm that outputs a CRT basis $c_1, \ldots, c_r \in \Lambda$ for those ideals.*

The proof is the same as in the ring case [27, Lemma 2.13]. Using Lemma 5 we can efficiently invert the natural CRT isomorphism. Given $a = (a_1, \ldots, a_r) \in \bigoplus_{i=1}^r (\Lambda/\mathcal{I}_i)$, it can be easily checked that its inverse is $b = \sum_{i=1}^r a_i c_i \mod \mathcal{I}$.

The next two lemmas will be required later to construct an efficiently invertible bijection between quotient spaces $\mathcal{I}/\langle q\rangle \cdot \mathcal{I}$ and $\Lambda/\langle q\rangle$.

Lemma 6. *Assume q is unramified in L. Let $\mathcal{I}$ be an ideal of the natural order Λ which is maximal and let $\mathcal{J} = q \cdot \Lambda = \langle q\rangle \cdot \Lambda$, where q is a prime integer and $\langle q\rangle = \prod_{i=1}^r \mathfrak{q}_i$ is a decomposition into prime ideals in $\mathcal{O}_K$. Assume $\gamma \notin \mathfrak{q}_i$ for each i. Then, there exists an element $t \in \mathcal{I} \cap \mathcal{O}_K$ such that the ideal $t \cdot \mathcal{I}^{-1} \subset \Lambda$ is coprime to $\mathcal{J}$, and we can compute such a t efficiently given $\mathcal{I}$ and the prime factorization of $\mathcal{J}$.*

Remark 1. The condition on γ will be immaterial in our use case, since when γ is a unit the only $\mathcal{O}_K$ ideal that contains γ is $\mathcal{O}_K$ itself. Meanwhile, the unramification of q will arise (relatively) naturally in the work, so it is not really a restriction.

Proof. For an ideal $\mathcal{I}$ denote by $\overline{\mathcal{I}}$ its intersection with K, which is a non-trivial ideal of $\mathcal{O}_K$ (see [33, Sect. 3]). We apply the corresponding [27, Lemma 2.14] to obtain $t \in \overline{\mathcal{I}}$ such that $t \cdot \overline{\mathcal{I}}^{-1}$ and $\overline{\mathcal{J}}$ are coprime as ideals of $\mathcal{O}_K$ and $t \in \overline{\mathcal{I}} \setminus \bigcup_{i=1}^r \mathfrak{q}_i \cdot \overline{\mathcal{I}}$. Assume, for a contradiction, that $t \cdot \mathcal{I}^{-1} + \mathcal{J} \neq \Lambda$ *i.e.,* the ideals are not coprime. Then, there is some maximal ideal $\mathcal{M}$ of Λ containing $t \cdot \mathcal{I}^{-1}$ and $\mathcal{J}$. Since q is unramified in L and $\gamma \notin \mathfrak{q}_i$, by [33, Propositions 1 and 4], this ideal must be one of the ideals $\mathfrak{q}_i \cdot \Lambda$ since it contains $\mathcal{J}$. Then $t \cdot \mathcal{I}^{-1} \subset \mathfrak{q}_i \cdot \Lambda$ and consequentially $t \in \mathfrak{q}_i \cdot \mathcal{I}$ because $\mathcal{I} \cdot \mathcal{I}^{-1} = \Lambda$ in a maximal order. Since t and $\mathfrak{q}_i$ are central, it follows that $t \in \mathfrak{q}_i \cdot \overline{\mathcal{I}}$, a contradiction. □

The next lemma will be the one we use in our reduction. As in RLWE, in practice we are interested in the case where $\mathcal{J} = \langle q\rangle$ for a prime integer q and $\mathcal{P} = \Lambda^\vee$. We will use the familiar notation $\mathcal{I}_q := \mathcal{I}/q \cdot \mathcal{I}$ for an ideal $\mathcal{I}$ and $q \in \mathbb{Z}$ throughout the paper.

Lemma 7. *Let Λ, γ and q be given in Lemma 6. Let $\mathcal{I}, \mathcal{J}$ be ideals of Λ, with $t \in \mathcal{I} \cap \mathcal{O}_K$ chosen as above such that $t \cdot \mathcal{I}^{-1}$ and $\mathcal{J}$ are coprime as ideals, and let $\mathcal{P}$ denote an arbitrary fractional ideal of Λ. Then, the function $\chi_t : \mathcal{A} \to \mathcal{A}$ defined as $\chi_t(x) = t \cdot x$ induces a module isomorphism from $\mathcal{P}/\mathcal{J} \cdot \mathcal{P} \to \mathcal{I} \cdot \mathcal{P}/\mathcal{I} \cdot \mathcal{J} \cdot \mathcal{P}$. Furthermore, in the case $\mathcal{J} = \langle q\rangle$ for a prime integer q we can efficiently compute the inverse.*

Proof. The proof is similar to that of [27]. Since t lies in the center of Λ, it is clear that multiplication by t induces a module homomorphism. Given the map $\chi_t : \mathcal{P} \to \mathcal{I} \cdot \mathcal{P}/\mathcal{I} \cdot \mathcal{J} \cdot \mathcal{P}$ and $j \in \mathcal{J} \cdot \mathcal{P}$, $\chi_t(j) = t \cdot j \in \mathcal{I} \cdot \mathcal{J} \cdot \mathcal{P}$, so it is clear that $\mathcal{J} \cdot \mathcal{P}$ is in the kernel of this map. Conversely, if $\chi_t(x) = 0$ then $t \cdot x \in \mathcal{I} \cdot \mathcal{J} \cdot \mathcal{P}$, from which it follows that $\mathcal{I}^{-1} \cdot t \cdot x \subset \mathcal{J} \cdot \mathcal{P}$. From the definition of coprime, $t \cdot \mathcal{I}^{-1} + \mathcal{J} = \Lambda$, from which it follows that there exists $a \in t \cdot \mathcal{I}^{-1}$, $b \in \mathcal{J}$ such that $a + b = 1$. Hence,

$x = (a+b) \cdot x = a \cdot x + b \cdot x$. Since $a \cdot x, b \cdot x \in \mathcal{J} \cdot \mathcal{P}$, it follows that $x \in \mathcal{J} \cdot \mathcal{P}$, from which injectivity follows immediately.

To demonstrate efficient invertibility, we must work slightly harder. Now let $\mathcal{J} = \langle q \rangle$. Compute t as in Lemma 6 and observe that the bijection $\chi_t : \Lambda_q \to \mathcal{I}_q$ is an additive homomorphism. Thus, it suffices to compute the inverse of all elements of a $\mathbb{Z}$ basis of $\mathcal{I}_q$, since then any element can be inverted by computing its representation in this basis and inverting that. We construct such a basis as follows. First, choose $n^2 \cdot d^4$ elements $x_i, i = 1, \ldots, n^2 \cdot d^4$ from Λ_q uniformly at random and compute $y_i = \chi_t(x_i)$ for each i. It follows that each y_i is a uniformly random element of $\mathcal{I}_q$. Then, with high probability the y_i's form a spanning set of $\mathcal{I}_q$ (see the proceeding lemma), which we can reduce to a $\mathbb{Z}$ basis $y'_1, \ldots, y'_{n \cdot d^2}$. This basis satisfies the desired property that each element has a known inverse. If this algorithm fails (e.g., there is no suitable basis $y'_1, \ldots y'_{n \cdot d^2}$), we repeat, choosing a fresh set of elements $x_1, \ldots, x_{n^2 \cdot d^4}$ until we succeed. □

Lemma 8. *Given a set of $n^2 \cdot d^4$ independent and uniformly random elements $\Xi \subset \mathbb{Z}_q^{n \cdot d^2}$, the probability that Ξ contains no set of $n \cdot d^2$ linearly independent vectors (over $\mathbb{Z}_q$) is exponentially small in d.*

This lemma is a straightforward adaptation of Corollary 3.16 of [42].

2.5. *Lattice Problems*

Computational problems on lattices represent the foundations of the security of (R)LWE, and will do so for our Cyclic LWE as well. The standard lattice problems are as follows.

Definition 11. Let $\| \cdot \|$ be some norm on $\mathbb{R}^n$ and let $\xi \geq 1$. Then the approximate Shortest Vector Problem (SVP_ξ) on input a lattice $\mathcal{L}$ is to find some nonzero vector $\mathbf{x}$ such that $\|\mathbf{x}\| \leq \xi \cdot \lambda_1(\mathcal{L})$.

Definition 12. Let $\| \cdot \|$ be some norm on $\mathbb{R}^n$ and let $\xi \geq 1$. Then the (approximate) Shortest Independent Vectors Problem (SIVP_ξ) on input a lattice $\mathcal{L}$ is to find n linearly independent nonzero vectors $\mathbf{x}_1, \ldots, \mathbf{x}_n$ such that $\max_i(\|\mathbf{x}_i\|) \leq \xi \cdot \lambda_n(\mathcal{L})$.

Definition 13. Let $\| \cdot \|$ be some norm on $\mathbb{R}^n$, let $\mathcal{L}$ be a lattice, and let $d < \lambda_1(\mathcal{L})/2$. Then the Bounded Distance Decoding problem ($\mathrm{BDD}_{\mathcal{L},d}$) on input $\mathbf{y} = \mathbf{x} + \mathbf{e}$ for $\mathbf{x} \in \mathcal{L}$ and $\|\mathbf{e}\| \leq d$ is to compute $\mathbf{x}$, or equivalently $\mathbf{e}$.

The above problems are all well investigated and believed to be sufficiently hard to base post-quantum cryptographic security on; there are no known algorithms for any of these problems (for suitable parameters) running in polynomial time in dimension n.

Unfortunately, these problems are not directly suitable for CLWE, where we will be interested in their adaptations to lattices generated by order ideals, similarly to how ideal lattices are used the ring case. Specifically we have the same problems on lattices that they induce under the map $\sigma_{\mathcal{A}}(\cdot)$. So, SVP becomes:

Definition 14. Let $\mathcal{A}$ be a cyclic algebra, let $\mathcal{I}$ be some (possibly fractional) ideal of the natural order Λ. Then, for an approximation factor $\xi \geq 1$, the $\mathcal{A}$-SVP$_\xi$ is to find a nonzero element $a \in \mathcal{I}$ such that $|a| := \|\sigma_{\mathcal{A}}(a)\|_2 \leq \xi \cdot \lambda_1(\mathcal{I})$, where as usual $\lambda_1(\mathcal{I})$ denotes the minimal length of nonzero elements of $\mathcal{I}$ in the given norm.

Remark 2. When we use these problems in our security reductions, we will assume that the ideals are in fact *integral* ideals (e.g., we exclude fractional ideals). Observe that this may be done without loss of generality, since solving the $\mathcal{A}$-SVP problem on the fractional ideal $\mathcal{I}$ may be done by solving it on the integral ideal $c\mathcal{I}$ (where $c \in K$ is the element such that $c\mathcal{I}$ is integral) and rescaling the solution.

Essentially we have a specialized version of the SVP problem; we must find an element of $\mathcal{I}$ with minimal norm (up to approximation factor) in the ideal $\mathcal{I}$. The extension of SIVP to $\mathcal{A}$-SIVP is analogous, but since we consider our objects as $\mathbb{Z}$-lattices we require the independent 'vectors' $a_1, \ldots, a_r$ to be linearly independent over $\mathbb{Z}$. For BDD, we need a suitable ambient space, and use the following definition.

Definition 15. Let $\mathcal{A}$ be a cyclic algebra, let $\mathcal{I}$ be some (possibly fractional) ideal of a maximal $\mathbb{Z}$-order Λ, and let $\delta < \lambda_1(\mathcal{I})/2$. Then the $\mathcal{A}$-BDD$_{\mathcal{I},\delta}$ problem, on input $y = x + e$ for $x \in \mathcal{I}$ and $e \in \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ satisfying $|e| \leq \delta$, is to compute x.

2.6. *The Learning with Errors Problem*

We will briefly recall the initial Learning With Errors (LWE) problem here; in Sect. 3 we will extend it to cyclic algebras. The problem comes in two forms; search and decision, both of which are based on the LWE distribution. Let n and q be positive integers, and let $\alpha > 0$ be some error parameter. Define $\mathbb{T} := \mathbb{R}/\mathbb{Z}$, the unit torus.

Definition 16. For a secret $\mathbf{s} \in \mathbb{Z}_q^n$, a sample $(\mathbf{a}, b) \leftarrow A_{\mathbf{s},\alpha}$ is taken by sampling a uniformly random vector $\mathbf{a} \in \mathbb{Z}_q^n$ and $e \leftarrow D_\alpha$ and outputting $(\mathbf{a}, b) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle / q + e \mod \mathbb{Z})$.

Given the above distribution, the LWE problem comes in two forms.

Definition 17. The search LWE problem is to recover $\mathbf{s}$ from a collection of samples $A_{\mathbf{s},\alpha}$. The decision LWE problem on input a collection of samples on $\mathbb{Z}_q^n \times \mathbb{T}$ is to decide whether they are uniform samples or were taken from $A_{\mathbf{s},\alpha}$ for some secret $\mathbf{s}$, where $\mathbf{s}$ is drawn uniformly at random from $\mathbb{Z}_q^n$.

Typically, the number of samples provided in each of these problems depends on the application. Since the decision problems has a probabilistic element, we will be interested in the advantage of the algorithms that solve it, which is defined as the difference between their acceptance probabilities on samples from an LWE distribution $A_{\mathbf{s},\alpha}$ and the uniform distribution. In practice, the decision problem is of more interest in cryptography.

We will not define the popular extensions of these problems to number fields or modules, known as Ring-LWE and Module-LWE, but the unfamiliar reader may find details in [27] and [22], respectively, both of which we reference frequently in this work.

3. The CLWE Problem

In this section we present the general definition of CLWE together with justifications for choices made in the definition, as well as constructions of specific algebras to use. We will save the security properties for Sect. 4.1.

Definition 18. Let L/K be a Galois extension of number fields of dimension $[L : K] = d$, $[K : \mathbb{Q}] = n$ with cyclic Galois group generated by $\theta(\cdot)$. Let $\mathcal{A} := (L/K, \theta, \gamma)$ be the resulting cyclic algebra with center K and invariant u with $u^d = \gamma \in \mathcal{O}_K$. Let Λ be an order of $\mathcal{A}$. For an error distribution ψ over $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$, an integer modulus $q \geq 2$, and a secret $s \in \Lambda_q^\vee$, a sample from the CLWE distribution $\Pi_{q,s,\psi}$ is obtained by sampling $a \leftarrow \Lambda_q$ uniformly at random, $e \leftarrow \psi$, and outputting $(a, b) = (a, (a \cdot s)/q + e \mod \Lambda^\vee) \in (\Lambda_q, \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}})/\Lambda^\vee$.

Remark 3. Unlike in commutative spaces, the order of multiplication of a and s is important; our choice is $(a \cdot s)$, but similar security properties would hold if one took $(s \cdot a)$ instead. Also observe that our modulo reduction in the second coordinate of the pair is well defined, since $(a \cdot s) \in \Lambda_q^\vee$.

As usual, the associated CLWE problem will come in search and decision variants.

Definition 19. Let Ψ be a family of error distributions over $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$. The search CLWE problem, which we denote by $\text{CLWE}_{q,s,\psi}$, is to recover s from a collection of independent samples from $\Pi_{q,s,\psi}$ for arbitrary $s \in \Lambda_q^\vee$ and $\psi \in \Psi$.

We do not state the number of samples allowed for this (or the next) problem, as typically it depends on the application.

Definition 20. Let Υ be some distribution on a family of error distributions over $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ and U_Λ denote the uniform distribution on $(\Lambda_q, (\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}})/\Lambda^\vee)$. Then, the decision CLWE problem, written $\text{D-CLWE}_{q,\Upsilon}$, is on input a collection of independent samples from either $\Pi_{q,s,\psi}$ for a random choice of $(s, \psi) \leftarrow U(\Lambda_q^\vee) \times \Upsilon$ or from U_Λ, to decide which is the case with non-negligible advantage.

3.1. *Discussions*

3.1.1. *Relation to Module-LWE*

First, we explain why we choose the order of multiplication $a \cdot s$. As discussed in the introduction, the transformation from a (primal) RLWE sample to n related LWE samples provides our motivation. Here, one RLWE sample $a \cdot s + e$, where $a, s, e \in R_q \cong \frac{\mathbb{Z}_q[x]}{x^n+1}$, generates n LWE samples by considering the multiplication operation as $A\mathbf{s} + \mathbf{e}$, where $A := \text{rot}(a)$ is a negacyclic matrix. For appropriate choices of error distributions, this is

precisely n LWE samples with the exception that there is some structure in the matrix A. By ordering the multiplication $a \cdot s$, we get a similar transform from CLWE to MLWE. Assuming for now that we have a discretized form of CLWE, and observing that for $q \in \mathbb{Z}$ we have $\Lambda_q \cong \bigoplus_{i=0}^{d-1} u^i \mathcal{O}_L/q\mathcal{O}_L$ (see [33]), we transform a CLWE sample $a \cdot s + e$ into matrix–vector form to get $\phi(a) \cdot \mathbf{s} + \mathbf{e}$, where $\mathbf{s}$ and $\mathbf{e}$ are vectors of dimension d over $\mathcal{O}_L/q\mathcal{O}_L$. Setting $A = \phi(a)$, one can see that for appropriate choices of error distribution this is similar to d samples from the MLWE distribution with some additional structure in the matrix A, as intended.

3.1.2. *The Natural Order vs. Maximal Order*

In this work we consider the case where the natural order Λ of $\mathcal{A}$ is also a maximal order. The benefit of using the natural order is that it is simple to construct and represent, whereas finding a maximal order is computationally slow. Additionally, the natural order is somewhat orthogonal, in the sense that it has the same span in each u^i coordinate independently of the other coordinates. This is advantageous when considering the relation to MLWE, where the module is always taken to be the full module $\mathcal{O}_K^d$.

As mentioned above, two-sided ideals in a maximal order form a free abelian group, which is not necessarily the case in the natural order. Further, as lattices, a maximal order gives denser (maximally so) sphere packing than the natural order, since the latter is a sublattice (of at least one maximal order). Fortunately, we will construct in Theorem 2 cyclic algebras whose natural order is also maximal, thus enjoying both the simplicity of the natural order and the convenience of a maximal order.

Example 2. Quaternion algebra over $\mathbb{Q}$ is defined by $\mathbb{H} = \{x + jy : x, y \in \mathbb{Q}(i)\}$, with the usual relations $i^2 = j^2 = -1$ and $ij = -ji$. It can be seen as a cyclic division algebra $(\mathbb{Q}(i)/\mathbb{Q}, \overline{(\cdot)}, -1)$ where $\overline{(\cdot)}$ denotes the complex conjugate and -1 is a non-norm element. A quaternion has matrix representation

$$\begin{pmatrix} x & -\overline{y} \\ y & \overline{x} \end{pmatrix}.$$

The *Lipschitz integers* $\mathcal{L} \subset \mathbb{H}$ form the (non-maximal) natural order $\mathcal{L} = \{x + jy : x, y \in \mathbb{Z}[i]\}$. The maximal Hurwitz order is given by

$$\mathcal{H} = \{a + bi + cj + d(-1 + i + j + ij)/2 : a, b, c, d \in \mathbb{Z}\}.$$

It is easy to check that, as $\mathbb{Z}$-lattices of dimension 4, the Lipschitz order is a sublattice of the Hurwitz order, of index 2.

3.1.3. *A Pair of Number Fields*

In MLWE, we are free to choose the dimension of our module over the underlying number field K. However, in the cyclic algebra case we are restricted to cases where we can find L, K, and γ such that $\mathcal{A} = (L/K, \theta, \gamma)$ is well defined. From a theoretical standpoint it is not immediately clear whether we want to consider asymptotic security in terms of n or d, but following our motivation from MLWE we suggest that n is likely

the suitable choice since the module dimension d is typically small in applications using MLWE, whereas the dimension of the underlying field K is large. However, there seems to be no a priori reason why with the right techniques one could not consider both n and d asymptotically; the only case a cyclic algebra precludes is high-dimensional MLWE over a low dimension number field L, because the parameter d occurs in both the module and field dimension.

3.2. *Evading BCV Style Attacks*

In our CLWE construction we have enforced that γ is selected so that $\mathcal{A}$ is a division algebra. We do this to avoid attacks in the style of [12] on the m-RLWE protocol. For $m = 2$, the m-RLWE protocol of [35] can be considered as a structured variant of MLWE, where the matrix A in the operation $A\mathbf{s} + \mathbf{e}$ is a negacyclic matrix over some ring R_q. More explicitly, 2-RLWE considers the tensor product of two fields $K = K_1 \otimes K_2$ and runs the LWE assumption in the ring of integers R_q. The example use case given in [35] considers power-of-two cyclotomics K_1, K_2 defined by the polynomials $x^{k_1} + 1$ and $y^{k_2} + 1$, respectively, claiming that the resulting problem in $R_q = \frac{\mathbb{Z}_q[x,y]}{(x^{k_1}+1, y^{k_2}+1)}$ effectively corresponds to an RLWE problem of dimension $k_1 \cdot k_2$ due to an obvious homomorphism between K and the two-power cyclotomic field L of degree $k_1 \cdot k_2$. The problem also represents a structured MLWE instance over $\frac{\mathbb{Z}_q[x]}{(x^{k_1}+1)}$ of dimension k_2.

However, the observation of [12] is that there is a smaller field K' containing K_1 such that there is a homomorphism from K into K' with a well-defined image for y. This is because the roots of distinct two-power cyclotomic polynomials are algebraically related. For example, in the case $k_1 = 8$, $k_2 = 4$, it is clear that the map taking y to x^2 and fixing K_1 is a well-defined homomorphism from K to K_1. Using this homomorphism, [12] simplifies the problem of solving one 2-RLWE instance by considering it as four RLWE instances in dimension k_1 rather than one instance in dimension $k_1 \cdot k_2$, essentially removing the module dimension k_2 from the problem.

We argue that the non-norm condition of γ precludes the existence of a homomorphism removing the module structure by taking a well-defined cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ to a smaller subfield containing K. We restrict our search to maximal subfields of $\mathcal{A}$, since any subfield is contained in at least one maximal subfield. It is a well-known result on division algebras that any maximal subfield E of $\mathcal{A}$ contains K and satisfies $[E : K] = d$, and that in the case of a cyclic division algebra $\mathcal{A}$ there is a choice of $u' \in \mathcal{A}$ such that the cyclic algebra $\mathcal{A}' := \bigoplus_j u'^j E$ is isomorphic to $\mathcal{A}$ (see Sect. 15.1, Proposition a of [41]). Assume, for a contradiction, that we had such a homomorphism $\chi : \mathcal{A} \to L$, where without loss of generality we assume the maximal subfield is L by the aforementioned proposition. Since L is Galois, the restriction of χ to L is an automorphism of L. It is clear that χ must agree on conjugates, since $\chi(u) \cdot \chi(\ell) = \chi(u \cdot \ell) = \chi(\theta(\ell) \cdot u) = \chi(u) \cdot \chi(\theta(\ell))$ for any $\ell \in L$. However, this contradicts χ being injective on L and it follows that no such homomorphism exists. Hence, we conclude that the attack style of [12] does not threaten our algebraic structure.

3.3. *Concrete Algebras for CLWE*

In order to apply the CLWE assumption in a practical cryptosystem, one must choose a concrete algebra as an ambient space. More generally, we are interested in finding families of algebras suitable for CLWE that allow for asymptotic analysis and varied security levels. Our search for algebras is motivated by the restrictions and conditions discussed in the previous section. In particular, we are interested in cyclic division algebras satisfying the following properties:

- The non-norm element γ must lie in $\mathcal{O}_K$ to keep the natural order closed under multiplication, and should satisfy $|\gamma| = 1$ in order to maintain both the coordinatewise independence and sub-multiplicative properties of the norm.[3]
- The dimension $n := [K : \mathbb{Q}]$ of the division algebra should be large and the degree $d := [L : K]$ should be small. This is to maintain the analogy with structured MLWE (the degree corresponds to the module rank) and follows from the search-decision reduction, which takes time polynomial in n but not in d.
- The base field K should be cyclotomic and q should split completely in K. This is also a result of the methodology of the search-decision reduction, which uses the well-understood factorization of $\langle q \rangle$ in $\mathcal{O}_K$. In addition, since the bulk of lattice-based cryptography is done over cyclotomic fields, we consider algebras which are small extensions of these as somewhat natural. We observe that an improved proof of decision security may allow this point to be dropped, whereas the other two points feel more integral.

Although significant effort has been expended by coding theorists to construct cyclic division algebras satisfying a variety of conditions, such as in [44] or [21], we find ourselves with a fairly unique set of restrictions. In particular, for reasons relating to desired applications, the majority of algebras used in coding theory are either of small total dimension or have small $[K : \mathbb{Q}]$ and scale asymptotically in $[L : K]$. Since we are interested in scaling up K asymptotically, we will have to build novel algebras satisfying the above requirements ourselves. We will, however, make heavy use of the following theorem as an intermediate step. Here ζ_m denotes a primitive mth root of unity where $\varphi(m) = n$ is the degree of the base field $K = \mathbb{Q}(\zeta_m)$.

Theorem 1. [21] *Let $m = p^a$ be a prime power and let $K = \mathbb{Q}(\zeta_m)$. Then, there exist infinitely many cyclic Galois extensions M/K of degree m such that ζ_m^i is not a norm of M/K for $0 < i < m$.*

We remark that the theorem is effective in the sense that it provides an explicit description of M, and we provide a summary of the recipe for constructing M. The crucial aspect of its construction is that M is a subfield of some cyclotomic extension of K, $K(\zeta_{q'})$ for a prime q', but we present its full description for completeness.

[3] We abbreviate the condition $|\sigma_i(\gamma)| = 1$ for all i by $|\gamma| = 1$, since in fact these are equivalent for algebraic γ.

First, find some prime q' such that $q' = 1 \mod p^a$ but $q' \neq 1 \mod p^{a+1}$, so that p^a is the highest power of p dividing $q' - 1$.[4] Set $M' = K(\zeta_{q'})$ so that by coprimality $M' = \mathbb{Q}(\zeta_{mq'})$. Then $\mathrm{Gal}(M'/K)$ is a cyclic group of order $q' - 1$ generated by some automorphism σ. Denote by M the subfield of M' fixed by σ^m. Then $[M : K] = m$ by the fundamental theorem of Galois theory and the extension is both cyclic and Galois. Finally, localization theory is used to show that the powers of ζ_m are not norms in this extension. In this way, the theorem constructs M explicitly.

The part of this theorem of our interest is that it allows us to scale K asymptotically, but this comes with a drawback of very high degree M, *i.e.*, it only permits a degree-m extension M of a degree-$\varphi(m)$ base field K. We present a new method that uses this theorem as a starting point to construct good algebras satisfying our restrictions. More precisely, our construction will begin with Theorem 1 and then use elementary methods from Galois theory to build more favorable fields.

3.3.1. *Constructions using Subfields*

We squash the field M from Theorem 1 to a subfield L of small index over the base K satisfying the necessary properties to generate a cyclic algebra.

Theorem 2. *Let $K = \mathbb{Q}(\zeta_m)$, where $\varphi(m) = n$, be a prime power cyclotomic with $m = p^a$ for some integer a and prime p. Then, there exists a cyclic Galois extension L/K of any index d dividing m within which ζ_m satisfies the non-norm condition.*

Remark 4. Since the proof will provide an explicit description of L, the correct interpretation of this theorem is that we can construct cyclic division algebras $\mathcal{A} = (L/K, \theta, \gamma)$ with $\langle\theta\rangle = \mathrm{Gal}(L/K)$, $\gamma = \zeta_m$, $K = \mathbb{Q}(\zeta_m)$, and $[L : K]$ is any divisor of $m = p^a$. Figure 2 shows all possible cases of intermediate field L between K and M.

Proof. Let $K = \mathbb{Q}(\zeta_m)$ for a fixed $m = p^a$ with prime p and integer a. Following the construction of Theorem 1 fix a cyclic Galois extension M/K of degree m such that ζ_m^i is not a norm of an element of M into K for any $i = 1, 2, \ldots, m - 1$. We will choose L as a suitable intermediate extension $M/L/K$. Let σ denote the generator of $\mathrm{Gal}(M/K)$, an automorphism of degree m. For d dividing m, σ^d fixes an extension L of K with $[M : L] = |\mathrm{Gal}(M/L)| = m/d$ and it follows from the tower lemma that $[L : K] = d$. We will show that L is a satisfactory extension of K.

First, since $\mathrm{Gal}(M/L)$ is a normal subgroup of $\mathrm{Gal}(M/K)$ we see that L/K is a normal, and hence Galois,[5] extension. It follows from standard Galois Theory that

$$\mathrm{Gal}(L/K) \cong \mathrm{Gal}(M/K)/\mathrm{Gal}(M/L).$$

Both groups in the quotient are cyclic, and so $\mathrm{Gal}(L/K)$ is cyclic with some generator θ. Furthermore, this isomorphism also allows us to deduce $|\mathrm{Gal}(L/K)| = d$.

[4] It is easy to show that infinitely many primes satisfying this condition always exist by appealing to classical theorems of Chebotarev or Dirichlet.

[5] Since in this case all extensions are separable.

Galois group fixing the field	Field	Degree over K
$\langle 1 \rangle$	M	p^a
$\langle \sigma^{p^{a-1}} \rangle$	L_{a-1}	p^{a-1}
	$\vdots$	
$\langle \sigma^{p^2} \rangle$	L_2	p^2
$\langle \sigma^p \rangle$	L_1	p
$\langle \sigma \rangle$	K	1

Fig. 2. Cyclic subfields between M and K from Galois correspondence. $\langle \sigma^i \rangle$ denotes the group generated by σ^i, where σ is the generator of $\mathrm{Gal}(M/K)$.

We've shown that L/K is a cyclic Galois extension of degree d; we are left to show that ζ_m^i is not a norm for $i = 1, \ldots, d-1$. Let $\overline{M}$ denote $N_{M/K}(M^\times)$ and $\overline{L}$ denote $N_{L/K}(L^\times)$. Say $\zeta_m^i \in \overline{L}$, fixing $x \in L$ such that $N_{L/K}(x) = \zeta_m^i$. Now by transitivity of the norm,

$$\begin{aligned} N_{M/K}(x) &= N_{L/K}(N_{M/L}(x)) \\ &= N_{L/K}(x^{m/d}) \\ &= \zeta_m^{(m/d)i} \end{aligned}$$

where the first equality follows from $x \in L$ and the second since the norm is multiplicative. $\overline{M}$ does not contain any power of ζ_m except $\zeta_m^m = 1$ since ζ_m is a non-norm element in M/K, so it follows that $m|(m/d)i$ and so $d|i$. From this we conclude that $\zeta_m, \zeta_m^2, \ldots, \zeta_m^{d-1}$ do not lie in $\overline{L}$ and so ζ_m satisfies the non-norm condition. □

Remark 5. We presented the proof in the above form for ease of legibility, but it is straightforward to extend the argument in the final paragraph to show that ζ_m^{jd+1} satisfies the non-norm condition for any $j = 0, 1, \ldots, (m/d) - 1$.

This is an effective construction that allows us to build cyclic division algebras of the form $\mathcal{A} = (L/K, \theta, \gamma)$ where $|\gamma| = 1$, K is an arbitrary prime power cyclotomic, and L is an extension of K with degree divisible by the prime p. For cryptographically relevant examples, we can consider degree 2 or 4 extensions of a 2-power cyclotomic or degree 3 extensions of a 3-power cyclotomic. Given the impossibility result of Appendix A and the restriction on the absolute value of γ, we view these algebras as essentially the best possible, at least for the case where K is a prime-power cyclotomic.

As discussed in Sect. 3.1, the natural order is not necessarily a maximal order. Nevertheless, the following theorem shows that the specific family of algebras we have constructed in Theorem 2 represents a lucky case (its proof is given in Appendix B).

Theorem 3. *For the family of cyclic division algebras $\mathcal{A} = (L/K, \theta, \zeta_m)$ constructed in Theorem 2, the natural order of $\mathcal{A}$ is maximal.*

This makes our constructed family of algebras very attractive, as it enjoys both the simplicity of the natural order and the nice property of a maximal order.

Remark 6. In the context of multiblock space-time coding [21], the construction of Theorem 1 allows for a space-time code for m antennas and $\varphi(m)$ blocks, *i.e.*, a relatively small number of blocks. With our new construction Theorem 2, any number $\varphi(mk)$, $k \in \mathbb{N}$ such that mk is a power of p, of blocks becomes possible. Further, using a maximal order leads to optimum coding gains; it was not realized in [21] that the natural order from Theorem 1 is actually maximal.

3.3.2. *Constructions using Compositum Fields*

The algebras with prime-power cyclotomic centers of the previous subsection use the field construction technique of Theorem 2, and as such they are restricted to algebras whose dimension N is in the form $p^k(p-1)$ for a prime p and integer k. We present another method of constructing algebras using compositum fields that allows us to target dimensions not achievable in this setting.

This method starts from extensions which are nearly what we are looking for and applies field compositums (cf. [43, Chapter 30]). Say we have a Galois field extension L'/K' with non-norm element $\gamma \in \mathcal{O}_{K'}$ whose Galois group is cyclic of degree d. Let F be some other Galois number field with $F \cap L' = \mathbb{Q}$. Then $\mathrm{Gal}(L'F/K'F) \cong \mathrm{Gal}(L'/K')$ and γ is a non-norm element in $L'F/K'F$. Relabeling this extension as L/K and letting θ denote the cyclic generator of the Galois group gives a cyclic field extension with non-norm γ such that $[L : K] = d$ and $[K : \mathbb{Q}] = [K' : \mathbb{Q}] \cdot [F : \mathbb{Q}]$. The relations among these fields are illustrated in Fig. 3a.

One can generalize this method to the case where the base field can not be written conveniently as a compositum of two fields. Let L'/K' be a cyclic Galois extension of degree d with non-norm element γ and let K be another Galois number field which contains K'. Then KL'/K is a cyclic Galois extension of degree k for some k dividing d, and in particular if $K \cap L' = K'$ then $k = d$ since the fields are linearly disjoint above K'. See Fig. 3b for the relations among these fields.

Similar to the subfield method, we also have the following theorem for the compositum field method (the proof is given in Appendix B).

Theorem 4. *Let $K = \mathbb{Q}(\zeta_n)$ where $n = p^r$ and p is prime, L/K be a finite cyclic extension of degree d with $Gal(L/K) = \langle\theta\rangle$ and $Gal(L/\mathbb{Q})$ abelian, and $F = \mathbb{Q}(\zeta_{q^t})$ where $F \cap L = \mathbb{Q}$. Suppose the natural order $\Lambda \subset \mathcal{A} = (L/K, \theta, \zeta_n)$ is maximal. Then, if $[F : \mathbb{Q}]$ and d are coprime, the natural order Λ' of the cyclic division algebra $\mathcal{A}' = (LF/KF, \theta', \zeta_n)$ is also maximal.*

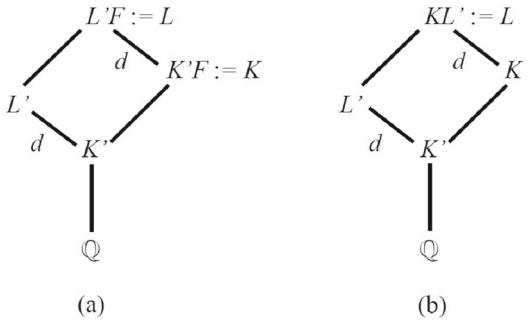

Fig. 3. Constructions using field compositums: **a** base field K is a compositum $K'F$, **b** K cannot be written as a compositum.

Table 1. Sample parameters of cyclic algebras.

Method	Center K	$n = [K : \mathbb{Q}]$	$d = [L : K]$	Total dimension $N = nd^2$ of $\mathcal{A}$
Subfield	$\mathbb{Q}(\zeta_{81})$	54	3	486
Subfield	$\mathbb{Q}(\zeta_{256})$	128	2	512
Subfield	$\mathbb{Q}(\zeta_{64})$	32	4	512
Subfield	$\mathbb{Q}(\zeta_{512})$	256	2	1024
Subfield	$\mathbb{Q}(\zeta_{128})$	64	4	1024
Subfield	$\mathbb{Q}(\zeta_{243})$	162	3	1458
Compositum	$\mathbb{Q}(\zeta_{192})$	64	3	576
Compositum	$\mathbb{Q}(\zeta_{576})$	192	2	768
Compositum	$\mathbb{Q}(\zeta_{384})$	128	3	1152

The subfield method is given in Sect. 3.3, while the compositum method is given in Appendix 3.3.2

3.4. *Sample Parameters*

Now that we have discussed our techniques for constructing suitable number fields we proceed to demonstrate that these methods are able to attain cryptographically relevant dimensions. In this section, we present a small selection of proof-of-concept dimensions in Table 1 where we take our motivation for choices of dimension from KYBER and NewHope, since they are the successful second round NIST candidates whose methods are most similar to our own. Thus, we aim for dimensions in the region of between 512 and 1024, dimensions proposed for both NewHope and KYBER (which also achieves dimension 768). Of course, these schemes are restricted to having power-of-two ring dimension n and so their choices of dimension may not be optimal in general, but FrodoKEM [13], a plain LWE scheme, suggests dimensions in around the same range, specifically 640, 976, and 1344, so we consider dimensions in this region a sensible starting point. Corresponding to KYBER and other MLWE-based schemes we will set a small 'module' rank $d := [\mathcal{A} : L]$. We are constricted in our choice of fields by the fact that d appears as a square in the total dimension $N = nd^2$, but for the most part we are able to work around this problem.

3.4.1. *Subfields*

Two-Power Cyclotomic K We begin with straightforward cases where we can apply Theorem 2 immediately to obtain fields in suitable dimensions. Let K be a two-power cyclotomic field, $K = \mathbb{Q}(\zeta_{2^k})$, with dimension $n := 2^{k-1}$. Since the rank $d = [L : K] = [\mathcal{A} : L]$ is a small power of two, the dimension n of K will be dictated by the choice of module rank d. We construct rank 2 and 4 examples as follows:

- For $d = 2$ we have $[\mathcal{A} : K] = 4$, so for total dimension 1024 we set $K = \mathbb{Q}(\zeta_{512})$.
- For $d = 4$ we have $[\mathcal{A} : K] = 16$, so for total dimension 1024 we set $K = \mathbb{Q}(\zeta_{128})$.

To obtain algebras in dimension 512, simply pick K with dimension $n/2$ e.g., $\mathbb{Q}(\zeta_{256})$ and $\mathbb{Q}(\zeta_{64})$, respectively. In all cases, Theorem 2 lets us pick the non-norm element γ as a root of unity.

Three-Power Cyclotomic K Since $3 \nmid 1024$, one cannot achieve algebras in dimension 1024 with a 3-power cyclotomic center and instead we set about searching for algebras of nearby dimensions. Although we are unable to build fields in this case with dimension around 1024, we can get close to the more lightweight cryptographic dimension of 512 used in schemes targeting a lower security level. Recall that if $K = \mathbb{Q}(\zeta_{3^k})$ then K has dimension $n := \phi(3^k) = 2 \cdot 3^{k-1}$. Again, the module rank is a power of 3 and the choice of module rank will define the choice of n.

- For $d = 3$ we have $[\mathcal{A} : K] = 9$, so for total dimension 486 we set $K = \mathbb{Q}(\zeta_{81})$. The next achievable dimension is 1458, for which $K = \mathbb{Q}(\zeta_{243})$.
- For $d = 9$ we have $[\mathcal{A} : K] = 81$. To achieve the same total dimensions we take small base fields $K = \mathbb{Q}(\zeta_9)$ and $\mathbb{Q}(\zeta_{27})$, respectively.

3.4.2. *Compositum Fields*

We give example algebras of dimensions 576, 768 and 1152 in Table 1 with less restrictive dimension using field compositum techniques. We propose two alternate methods of applying field compositums in Fig. 3a: either use Theorem 2 to make an algebra which already has large dimension by selecting large center K and small extension L, then compose a small field F onto K and L to tweak the total dimension. Alternatively, one can create algebras by selecting small fields L and K using Theorem 1 and composing both with a large field F.

We begin with an example of the first method that achieves dimension 768. Let L' be a degree two extension of the field $K' = \mathbb{Q}(\zeta_{64})$ chosen by Theorem 2 with non-norm root of unity γ, so that the corresponding algebra $\mathcal{A}'$ has dimension 128. Compose both L' and K' with the field $F = \mathbb{Q}(\zeta_9)$, denoting the compositums by L and K respectively. Then γ is still a non-norm element in the extension L/K, a degree two extension that is cyclic and Galois, and the algebra $\mathcal{A} = (L/K, \theta, \gamma)$ is a cyclic algebra of dimension $6 \times 128 = 768$, as required. We observe that here the center K corresponds to the fields with fast operations used in [29].

Our final method of composing large degree fields onto small degree extensions is aimed at targeting odd module ranks. Begin by choosing the desired module rank d as a (likely small) odd prime. Then set $K' = \mathbb{Q}(\zeta_d)$ and pick L' as a cyclic Galois extension of K' in which the dth root of unity is a non-norm element using Theorem 1. Let $F := \mathbb{Q}(\zeta_{2^k})$ and again let L and K denote its compositum with L' and K' respectively.

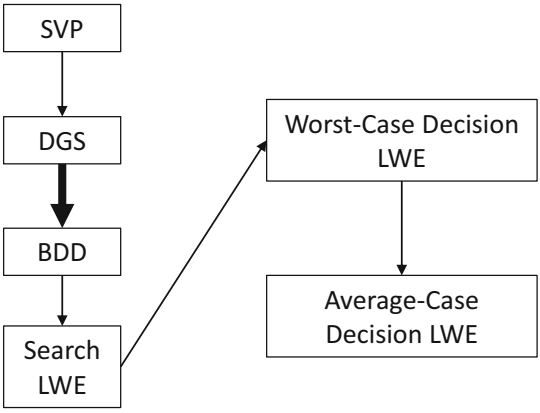

Fig. 4. Reductions for LWE. The bold arrow denotes a quantum step.

Then $\mathcal{A} = (L/K, \theta, \gamma)$ is a cyclic algebra with $n := [K : \mathbb{Q}] = (d-1)2^{k-1}$ and $d = [L : K]$ a small prime. The form of the total dimension $N = d^2(d-1)2^{k-1}$ constrains our choice of dimension, but for examples of cryptographically relevant sizes with $d = 3$ one can consider setting $k = 6$ or $k = 7$ to achieve dimension $N = 576$ or $N = 1152$ respectively.

3.5. *Extensions where q Splits Completely*

All suggested algebras in the previous section satisfy the conditions required for our chosen norm $\|\sigma_{\mathcal{A}}(x)\|_2$ to be well defined. In particular, they have root of unity non-norm γ and K is cyclotomic. Because any $q = 1 \mod m$ splits completely in $\mathbb{Q}(\zeta_m)$, it is straightforward to find q which splits completely in $\mathcal{O}_K$.

Later in this paper, in order to enable efficient multiplication algorithms, it will turn out that it is convenient to have a modulus q that splits completely into a product of prime ideals in both $\mathcal{O}_K$ and $\mathcal{O}_L$. Recall Lemmas 6 and 7 also require q be unramified in L. An appeal to Chebotarev's Density Theorem suggests that a proportion of $1/d$ of the primes q that split completely in K also do so in L. In cases where d is small this suggests that finding such primes should not prove too arduous; but since cryptosystems require specific parameters rather than density arguments, we provide constructions satisfying the requisite conditions on q in Appendix C.

4. Security Proof

The 'standard' security reductions used in [42] and [27] firstly reduce certain lattice problems to search LWE and RLWE, then establish hardness of the decision problem via a search-decision reduction. This proof follows a sequence of shorter reductions as shown in Fig. 4.

The reduction from the approximate SVP to the search LWE problem implies that search LWE is at least as hard as approximate SVP. It can be explained as follows:

first, the approximate SVP is reduced to the problem of sampling a discrete Gaussian of narrow variance over a lattice, where intuitively sampling from a sufficiently narrow Gaussian should output a vector whose norm is reasonably short compared to the first minima. Then, a quantum algorithm reduces the problem of sampling from a narrow Gaussian to that of solving the BDD problem on the dual lattice. Finally, a transformation maps an instance of the BDD problem to an appropriate instance of the LWE problem, reducing the BDD problem to that of search LWE.

For applications in cryptography, the hardness of the decision problem is preferred to that of the search problem. Assuming that the decision problem is hard implies that LWE samples are computationally indistinguishable from uniform, so intuitively an LWE sample can be used to hide a message m as an element of $\mathbb{Z}_q^n$ by adding it to b.

Using similar machinery, we reduce a BDD problem to search CLWE using the same method as in [27]. The methodology of their search-decision reduction is an adaptation of that of Regev's, which relies on guessing each coordinate of the secret $\mathbf{s}$ separately. The adaptation to the ring case instead guesses the coordinate of the secret ring element s modulo a suitable collection of ideals $\mathfrak{p}_i$ such that guessing $s \mod \mathfrak{p}_i\mathcal{O}_K^\vee$ requires only a polynomial number of guesses, from which s is recovered using the CRT. We apply a similar method in suitable subrings to deduce the hardness of our decision problem. The main technical novelty is to deal with non-commutativity in the proof.

For the remainder of this paper, we will always be working in an extension of number fields L/K, where $[L : \mathbb{Q}] = [L : K] \cdot [K : \mathbb{Q}] = d \cdot n$. Recall from the motivation of structured MLWE and the sample algebras given that in practice we seek asymptotic security in n, since the parameter d corresponds to the typically small module dimension.

4.1. *Hardness of Search CLWE*

In the following, let $\mathcal{A}$ be a cyclic division algebra over a number field L with center K and natural, maximal order Λ with $|\gamma| = 1$. Let $\alpha = \alpha(n) \in (0, 1)$ and $q = q(n) \geq 2$, unramified in L, be parameters such that $\alpha \cdot q \geq \omega(\sqrt{\log N})$. We denote by $\mathcal{A}-\mathrm{DGS}_\xi$ the problem of sampling a discrete Gaussian $D_{\mathcal{I},\xi}$, where $\mathcal{I}$ is some ideal of the order Λ. Also denote by N the total dimension of $\mathcal{A}$, $N := nd^2$.

For the reduction of BDD to Search CLWE, we begin with the cyclic algebra analogy of the BDD-to-LWE samples transformation from Sect. 4 of [27]. As is standard for LWE security, we use the following 'modulo q' definition of BDD:

Definition 21. For any $q \geq 2$ the $q\mathcal{A}-\mathrm{BDD}_{\mathcal{I},\delta}$ problem is as follows: given an instance of the $\mathcal{A}-\mathrm{BDD}_{\mathcal{I},\delta}$ problem $y = x + e$ with solution $x \in \mathcal{I}$ and error $e \in \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ satisfying $\|e\|_{2,\infty} \leq \delta$, output $x \mod q\mathcal{I}$.

We use (a special case of) Lemma 3.5 from [42], which lifts immediately since it is lattice preserving.

Lemma 9. *For any $q \geq 2$ there is a deterministic polynomial time reduction from $\mathcal{A}-BDD_{\mathcal{I},\delta}$ to $q\mathcal{A}-BDD_{\mathcal{I},\delta}$.*

We now present an algorithm which transforms $q\mathcal{A}$-BDD samples to CLWE samples given some additional Gaussian samples. The algorithm is the same in spirit as Lemma 4.7 of [27], but has some technical differences induced by the structure of cyclic algebras.

Lemma 10. *Let $\mathcal{A}$ be as above. There is a probabilistic polynomial time algorithm that on input a prime integer $q \geq 2$, a fractional ideal $\mathcal{I}^\vee \subset \Lambda$, a $q\mathcal{A}-BDD_{\mathcal{I}^\vee,\alpha q\cdot\omega(\sqrt{\log(nd)})/\sqrt{2nd}\cdot r}$ instance $y = x + e$ where $x \in \mathcal{I}^\vee$, a parameter $r \geq \sqrt{2}q \cdot \eta(\mathcal{I})$, and samples from the discrete Gaussian $D_{\mathcal{I},r'}$ with $r' \geq r$, outputs samples that are within negligible statistical distance of the CLWE distribution $\Pi_{q,s,\Sigma}$ for a secret $s = \chi_t(x \mod q\mathcal{I}^\vee) \in \Lambda_q^\vee$, where χ_t is as in Lemma 7 and Σ is an error distribution such that in the case where $|\gamma| = 1$ the resulting error e'' has marginal distribution in its i, jth coordinate that is Gaussian with parameter $r_{i,j} \leq \alpha$.*

Proof. The proof will be in two parts—first, we will describe the algorithm, then we will prove correctness.

Begin by computing an element $t \in \mathcal{I}$ such that $\mathcal{I}^{-1} \cdot \langle t \rangle$ and $\langle q \rangle$ are coprime using Lemma 6. We can now create a sample from the CLWE distribution as follows: take an element $z \leftarrow D_{\mathcal{I},r'}$ from the Gaussian samples, and compute a pair

$$(a, b) = (\chi_t^{-1}(z \mod q\mathcal{I}), (z \cdot y)/q + e' \mod \Lambda^\vee) \in (\Lambda_q \times (\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}})/\Lambda^\vee)$$

where $e' \leftarrow D_{\alpha/\sqrt{2}}$.

We now claim that these samples are within negligible statistical distance of the CLWE distribution and that s is uniformly random. First we show that $a \in \Lambda_q$ is statistically close to uniform. By assumption, $r \geq q \cdot \eta(\mathcal{I})$ and so by appealing to Lemma 1 it can be seen that any value $z \mod q\mathcal{I}$ is obtained with probability in the interval $[\frac{1-\varepsilon}{1+\varepsilon}, 1] \cdot \beta$ for some positive β, from which it follows immediately that the statistical distance between $z \mod q\mathcal{I}$ and the uniform distribution is bounded above by 2ε. Since χ_t of Lemma 7 and its inverse are both bijections, we conclude that $a = \chi_t^{-1}(z \mod q\mathcal{I})$ is within statistical distance 2ε of the uniform distribution over Λ_q.

Now we must show that b is in the form $(a \cdot s)/q + e''$, for some suitable error e'' and a uniformly random s, where we condition on some fixed value of a. By construction,

$$\begin{aligned} b &:= (z \cdot y)/q + e' \mod \Lambda^\vee \\ &= (z \cdot x)/q + (z \cdot e)/q + e' \mod \Lambda^\vee, \end{aligned}$$

so since $z = t \cdot a \mod \Lambda_q^\vee$ and t lies in the center of $\mathcal{A}$ it follows that $(z \cdot x)/q = (a \cdot t \cdot x)/q = (a \cdot s)/q \mod \Lambda^\vee$ for $s := \chi_t(x \mod q\mathcal{I}^\vee)$. It follows that s is uniformly random over $\Lambda_q^\vee$ as long as x is uniform over $\mathcal{I}^\vee$, since χ_t is a bijection.

Finally it is left to show that, conditioned on a fixed value of a, the marginal distribution of the i, jth coordinate of the error term $e'' = (z \cdot e)/q + e'$ is negligibly close to that

specified by Σ. We can explicitly calculate the error as

$$e'' = \sum_{i=0}^{d-1} u^i \left(\sum_{j+k=i} \theta^k(z_j) \cdot e_k(1-(1-\gamma)\mathbb{1}_{j+k\geq d}) \right) + e' \tag{1}$$

where the sum $j+k$ is taken modulo d and the function $(1-(1-\gamma)\mathbb{1}_{j+k\geq d})$ is 1 if $j+k<d$ and γ otherwise.[6] Since $|\gamma|=1$ and $z \leftarrow D_{\mathcal{I},r}$ is spherically distributed, it follows that multiplying by γ and applying the permutation of j coordinates induced by θ does not change the distribution of $z_{i,j}$. Hence, each marginal distribution may be analyzed independently as in the case of MLWE, and the result follows using the analysis of the error from Lemma 4.15 of [22]. □

Though we do not specify the covariance of Σ, one can see that each entry of $\sigma_{\mathcal{A}}(z)$ appears in $\sigma_{\mathcal{A}}(e'')$ exactly d times, and so by symmetry each element of $\sigma_{\mathcal{A}}(e'')$ has nonzero correlation with at most d^2 other entries. Hence, a proportion of at most $\frac{nd^4}{n^2d^4}=\frac{1}{n}$ of entries of Σ are nonzero. This is the family of error distributions we will claim hardness of search CLWE for; we remark that it is a Gaussian distribution whose marginals are Gaussian with variance at most α.

Definition 22. We define the family of error distributions Σ_α as the set of all Gaussian distributions Σ over $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ whose marginal distribution in its (i,j)th coordinate is Gaussian with parameter $r_{i,j} \leq \alpha$.

The following theorem is our analogy of Lemma 4.10 of [22].

Theorem 5. *Given an oracle that solves $CLWE_{q,\Sigma_\alpha}$ for input $\alpha \in (0,1)$, an integer $q \geq 2$, an ideal $\mathcal{I} \subset \Lambda$, a number $r \geq \sqrt{2}q \cdot \eta_\varepsilon(\mathcal{I})$ satisfying $r' := r \cdot \omega(\sqrt{\log N})/(\alpha q) > \sqrt{2N}/\lambda_1(\mathcal{I}^\vee)$, and polynomially many samples from the discrete Gaussian $D_{\mathcal{I},r}$ there exists an efficient quantum algorithm that outputs an independent sample from $D_{\mathcal{I},r'}$.*

As usual, we obtain Theorem 5 in two steps, first the main reduction of Lemma 10, then the following quantum step adapted from [42]. We use a form of $\mathcal{A}-\mathrm{BDD}_{\mathcal{I},\delta}$ from [22] where we bound the offset in the norm $\|e\|_{2,\infty} := \max_j \sqrt{(\sum_{i=0}^{d-1} |\sigma_j(e_i)|^2)} \leq \delta$, where σ denotes the canonical embedding of $\mathcal{I}$.

Lemma 11. *There is an efficient quantum algorithm that given any $N = n \cdot d^2$-dimensional lattice from some ideal $\mathcal{I}$, a real $\delta < \lambda_1(\mathcal{I}^\vee)/(2\sqrt{2nd})$, and an oracle that solves $\mathcal{A}$-$BDD_{\mathcal{I}^\vee,\delta}$ with all but negligible probability, outputs an independent sample from $D_{\mathcal{I},\sqrt{d}\omega(\sqrt{\log(nd)})/\sqrt{2}\delta}$.*

We can then prove Theorem 6 in the standard iterative manner; for a very large value of r, e.g., $r \geq 2^{2N}\lambda_N(\mathcal{I})$, start by sampling classically from $D_{\mathcal{I},r}$. Then apply

[6] This term is just indicating whether or not we have had to use the relation $u^d = \gamma$ in this summand or not.

the above algorithm to obtain a polynomial number of samples from $D_{\mathcal{I},r'}$. Repeating this step gives samples from progressively narrower distributions, until we arrive at the desired Gaussian parameter $s \geq \xi$. In order to classically sample the initial collection of Gaussian samples, we use the standard Lemma 3.2 of [42] to sample $D_{\mathcal{I},r}$ on the module representation $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$.

Theorem 6. *Let $\mathcal{A}$ be a cyclic division algebra over a number field L with center K and natural, maximal order Λ with $|\gamma| = 1$. Let $\alpha = \alpha(n) \in (0,1)$ and $q = q(n) \geq 2$, unramified in L, be parameters such that $\alpha \cdot q \geq \omega(\sqrt{\log N})$. Then, there is a polynomial-time quantum reduction from $\mathcal{A}$-DGS$_\xi$ to search CLWE$_{q,\Sigma_\alpha}$ for any $\xi = r \cdot \sqrt{d}\omega(\sqrt{\log(d \cdot n)})/\alpha q$, where $r > \sqrt{2}q \cdot \eta_\varepsilon(\mathcal{I})$.*

From this we deduce the following corollary, similarly to [22], since the lattice structure of our algebra is merely a special case of their modules.

Corollary 1. *Let $\mathcal{A}$, Λ, α and q be as above. Then, there is a polynomial-time quantum reduction from $\mathcal{A}$-SIVP$_\xi$ to search CLWE$_{q,\Sigma_\alpha}$ for any $\sqrt{8Nd} \cdot \xi = (\omega(\sqrt{dn})/\alpha)$.*

4.2. *Search To Decision Reduction*

In this subsection we will show that the hardness of decision CLWE follows from that of the search problem. Once again, we will follow a combination of the expositions of [27] and [22] for the ring and module cases, making necessary changes for the structure of cyclic algebras. We will make heavy use of the following CRT style decomposition, a rephrasing of [33, Lemma 4].

Lemma 12. *Let Λ be the natural order of a cyclic division algebra $\mathcal{A} = (L/K, \theta, \gamma)$ and let $\mathcal{I}$ be an ideal of $\mathcal{O}_K$ which splits completely as $\mathcal{I} = \mathfrak{q}_1 ... \mathfrak{q}_n$ as an ideal of $\mathcal{O}_K$. Then, we have the isomorphism*

$$\Lambda/\mathcal{I}\Lambda \cong \mathcal{R}_1 \times ... \times \mathcal{R}_n,$$

where $\mathcal{R}_i = \bigoplus_{j=0}^{d-1} u^j(\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L)$ is the ring subject to the relations $(\ell + \mathfrak{q}_i\mathcal{O}_L)u = u(\theta(\ell) + \mathfrak{q}_i\mathcal{O}_L)$ and $u^d = \gamma + \mathfrak{q}_i$.

Of course, this is not a true CRT decomposition, because we are considering ideals of $\mathcal{O}_K$ rather than those of Λ. In the case where γ is a unit, $\Lambda^\vee = \bigoplus_i u^i \mathcal{O}_L^\vee$ and the above lemma is also valid in the case where each instance of $\mathcal{O}_L$ and Λ are replaced with their respective duals. Also note that γ is a non-norm element in this lemma. The reduction from DGS to search CLWE requires Λ to be maximal, and currently the only known value of γ which makes the natural order maximal is an n-th root of unity, which is also a non-norm element. So these conditions are consistent.

As in [27], our reduction will be limited to certain choices of algebras. The above lemma considers the splitting of the ideal $\mathcal{I}$ as an ideal of the base field K. Setting $\mathcal{I} = \langle q \rangle$, the ideal generated by the modulus q, we will consider cases where q splits

completely in the base field. Now consider the family of algebras $\mathcal{A}$ in Sect. 3.3 and let $K = \mathbb{Q}(\zeta_{p^a})$ have dimension n. It follows that if $q \equiv 1 \mod p^a$ then q splits completely into a product of prime ideals $\mathfrak{q}_1, \ldots, \mathfrak{q}_n$ as an ideal of $\mathcal{O}_K$. Hence, we obtain the decomposition

$$\Lambda/q\Lambda \cong R_1 \times \ldots \times R_n$$

where R_i is as in Lemma 12.

Also as in [27], we see no way to avoid randomizing the error distribution in the resulting decision problem. Further, we show that an oracle for D-CLWE$_{q,\Upsilon_\alpha}$ on an algebra $\mathcal{A} = (L/K, \theta, \gamma)$ is also an oracle for the decision problem on any algebra $\mathcal{A}' = (L/K, \theta, \gamma')$ over the same number fields L, K and some other root of unity $\gamma' \in \mathcal{O}_K$. Intuitively this implies that for fixed L and K as in Sect. 3.3 the hardness of the D-CLWE problem is invariant under the choice of root of unity γ and will be required for Lemma 15. This is because there exist efficient, easy-to-compute isomorphisms sending $\mathcal{A}$ to $\mathcal{A}'$, which we will define shortly. The security reduction is similar in spirit to that for Ring LWE applying field automorphisms.

The main theorem of this subsection is Theorem 7 (given in the end of this subsection); we emphasize that our algorithm is only intended to be efficient in the dimension n of the base field K, since we expect to fix d as a small constant in practice. We will prove Theorem 7 in the usual manner: first we show that it is sufficient to recover the value of $s \in \Lambda^\vee/q\Lambda^\vee$ in one of the rings R_i (Lemma 13). Then, we use a hybrid distribution to define a decision problem in R_i, for which we demonstrate a search to decision reduction (Lemma 14). We then use a hybrid argument to conclude the proof (Lemma 16).

4.2.1. *CLWE in R_i*

In this section we will abuse notation and denote by $s \mod R_i$ the value of $s \in \Lambda^\vee/q\Lambda^\vee$ in the R_i coordinate under the isomorphism of Lemma 12.

Definition 23. The $R_i - \text{CLWE}_{q,\Sigma_\alpha}$ problem is to find the value $s \mod R_i$ given access to the CLWE distribution $\Pi_{q,s,\Sigma}$ for some arbitrary $\Sigma \in \Sigma_\alpha$.

In the following lemmata we make use of the automorphisms of K coordinatewise on the rings R_i. Since K is a Galois extension of $\mathbb{Q}$ and q splits completely, it follows that the automorphisms σ_i of K act transitively on the ideals $\mathfrak{q}_i$. We demonstrate how to extend these to functions of $\mathcal{A}$. First, extend these automorphisms to automorphisms α_i of L in some arbitrary manner. Then, we can extend these to isomorphisms $\alpha_i : \mathcal{A} \to \mathcal{A}'$, with $\mathcal{A}' = (L/K, \theta, \gamma')$, which agree with α_i on L and send u to u' with $u'^d = \alpha_i(\gamma)$ and $xu' = u'\theta(x)$ for $x \in L$. By the construction of K from [21], $\alpha_i(\gamma)$ is a non-norm element since it is some primitive nth root of unity, and so it is easy to check that this $\mathcal{A}'$ is a well-defined division algebra and that α_i is indeed an isomorphism which sends $\mathcal{A}$ to $\mathcal{A}'$. Furthermore, it fixes the family of error distributions Σ_α. This is because each component of $z \cdot e + e'$ is defined coordinatewise over the d copies of $L_\mathbb{R}$ in the module representation of $\mathcal{A}$, and since α_i induces the same permutation of the entries of the canonical embedding of L in each coordinate as an automorphism of L it fixes the family of choices for each of z, e, e'; hence, since α_i is an isomorphism the family

of distributions $z \cdot e + e'$ is fixed. It follows that the extended α_i function maps the $R_i-\mathrm{CLWE}_{q,\Sigma_\alpha}$ problem in $\mathcal{A}$ to the same problem in $\mathcal{A}'$, and moreover that this map preserves $\Lambda^\vee$ and the CRT style decomposition (Lemma 12) of $\Lambda_q^\vee$ by sending R_i to some R_j, where j depends on the choice of σ_i. We are now ready for the first step of our reduction.

Lemma 13. *There is a deterministic polynomial time reduction from $CLWE_{q,\Sigma_\alpha}$ to $R_i CLWE_{q,\Sigma_\alpha}$.*

Proof. Let $\mathcal{O}_i$ be an oracle for the $R_i - \mathrm{CLWE}_{q,\Sigma_\alpha}$ problem. Since Lemma 12 defines an isomorphism, it is sufficient to use $\mathcal{O}_i$ to solve the $R_j-\mathrm{CLWE}_{q,\Sigma_\alpha}$ for each j. Let $\alpha_{j/i}$ be an extension of the automorphism of K mapping $\mathfrak{q}_j$ to $\mathfrak{q}_i$, which exists by transitivity. Then, given a sample $(a, b) \leftarrow \Pi_{q,s,\Sigma}$, we construct the sample $(\alpha_{j/i}(a), \alpha_{j/i}(b))$. Since Λ_q and $\Lambda_q^\vee$ are fixed by each $\alpha_{j/i}$, the resulting pair is a valid CLWE sample in $\mathcal{A}' = (L/K, \theta, \alpha_{j/i}(\gamma))$; feeding these samples into $\mathcal{O}_i$ outputs a value $t_j \mod R_i$.

We claim $\alpha_{j/i}^{-1}(t_j) = s \mod R_j$. Since $\alpha_{j/i}$ is an automorphism, each sample (a, b) is mapped to a new CLWE sample $(\alpha_{j/i}(a), \alpha_{j/i}(a \cdot s/q + e) \mod \Lambda^\vee)$ in a new algebra $\mathcal{A}'$. We may write the second coordinate as $\alpha_{j/i}(a) \cdot \alpha_{j/i}(s)/q + \alpha_{j/i}(e) \mod \Lambda^\vee$. Since our automorphisms fix our family of error distributions Σ_α and map the uniform distribution to the uniform distribution, it follows that this is a valid CLWE instance with secret $\alpha_{j/i}(s)$ and error distribution $\Sigma' \in \Sigma_\alpha$. Hence, $\mathcal{O}_i$ outputs $t = \alpha_{j/i}(s) \mod R_i$, from which we recover $\alpha_{j/i}^{-1}(t) = s \mod R_j$, as required. □

4.2.2. *Hybrid CLWE and Search-Decision*

For this section we must introduce the cyclic algebra analog of the Hybrid LWE distribution used in [27]; we use the decomposition into the rings R_i rather than the CRT.

Definition 24. For a secret $s \in \Lambda_q^\vee$, distribution Σ over $\bigoplus_j u^j L_{\mathbb{R}}$, and $i \in [n]$, we define a sample from the distribution $\Pi_{q,s,\Sigma}^i$ over $\Lambda_q \times (\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}})/\Lambda^\vee$ by taking $(a, b) \leftarrow \Pi_{q,s,\Sigma}$ and $h \in \Lambda_q^\vee$ which is uniformly random and independent $\mod R_j$, $j \le i$ and $0 \mod R_j$, $j > i$, and outputting $(a, b + h/q)$. If $i = 0$, we define $\Pi_{q,s,\Sigma}^0 = \Pi_{q,s,\Sigma}$.

Using this distribution we define a worst-case decision problem relative to one R_i and reduce it to the search problem $R_i-\mathrm{CLWE}$.

Definition 25. For $i \in [n]$ and a family of distributions Σ_α, the W-D-$\mathrm{CLWE}_{q,\Sigma_\alpha}^i$ problem is defined as the problem of finding j given access to $\Pi_{q,s,\Sigma}^j$ for $j \in \{i-1, i\}$ and valid CLWE secret s and error distribution $\Sigma \in \Sigma_\alpha$.

For a technical reason in the following proof, we restrict our secret s so that $s \mod \mathcal{R}_i$ lies in a set G_i with the property that $g \neq h \in G_i$ implies $g - h$ is an invertible element. Applying this restriction for each i places $s \in G$ for a set $G = G_1 \times \cdots \times G_n$ of size $|G| = \prod_i |G_i|$. We will call such a set G a *pairwise different set*. We need to guarantee

that there exist sufficiently large choices of G. It is not difficult to see that the maximal set sizes $|G_i| = q^d$ and $|G| = q^{nd}$, because any set of matrices in $M_{d\times d}(\mathbb{F}_q)$ of size at least $q^d + 1$ contains two matrices with the same first row, whose difference is therefore uninvertible. Constructions of such maximal sets G are given in Appendix D.

Lemma 14. *Assuming constant d and $s \in G$, there is a probabilistic polynomial-time reduction from R_i−CLWE$_{q,s,\Sigma_\alpha}$ to W-D-CLWE$^i_{q,\Sigma_\alpha}$ for any $i \in [n]$.*

Proof. We follow the standard search-decision methodology of guessing the value of the secret mod R_i and then modifying the samples so that the decision oracle tells us whether or not our guess was correct. Note that there are only $|G_i|$ possible values of $s \mod R_i$, which is bounded above by q^{d^2}, polynomial in n, and so we may efficiently enumerate over the possible values.

We define the transform which takes a value $g \in \Lambda_q^\vee$ and maps $\Pi_{q,s,\Sigma}$ to $\Pi^{i-1}_{q,s,\Sigma}$ if $g = s \mod R_i$ or $\Pi^i_{q,s,\Sigma}$ otherwise as follows. On input a CLWE sample $(a, b) \leftarrow \Pi_{q,s,\Sigma}$, output the pair

$$(a', b') = (a + v, b + (h + vg)/q) \in \Lambda_q \times (\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}})/\Lambda^\vee,$$

where $v \in \Lambda_q$ is uniformly random mod R_i and $0 \mod R_j$ for $j \neq i$ and $h \in \Lambda_q^\vee$ is uniformly random and independent mod R_j, $j < i$ and 0 on the other R_j. It is clear that a' is still uniformly distributed on Λ_q, so we are left to show b' is correctly distributed. For a fixed value of a', we write

$$\begin{aligned} b' &= b + (h + vg)/q \\ &= (as + h + vg)/q + e \\ &= (a's + h + v(g - s))/q + e, \end{aligned}$$

where e is still drawn from Σ. If $g = s \mod R_i$, then $v(g - s) = 0 \mod R_i$, and so the distribution of the pair (a', b') is precisely $\Pi^{i-1}_{q,s,\Sigma}$. Otherwise, $v(g - s)$ is uniformly random mod R_i by assumption on G and 0 mod the other R_j, and so letting $h' = h + v(g - s)$ we see that the distribution of (a', b') is precisely $\Pi^i_{q,s,\Sigma}$. □

Remark 7. This is the only stage of the proof which enforces that the asymptotic complexity scales only with n and not with d, since we are forced to guess all of s mod R_i at once.

Since the above reduction is secret preserving, the required decision oracle for W-D-CLWE$^i_{q,\Sigma_\alpha}$ has the additional restriction that $s \in G$, but for the purposes of the rest of our proof it will be more convenient to have access to an oracle solving the at least as hard problem where s is arbitrary. Additionally, in practical applications we will use the decision problem for arbitrary s, so we see no benefit of the tighter reduction where s is restricted.

4.2.3. *Worst-Case to Average-Case Decision Reduction*

Now that we have removed the restriction that $s \in G$, we are able to follow the skeleton of the RLWE search-decision reduction of [27] more liberally.

Definition 26. The error distribution Υ_α on the family of possible error distributions is sampled from by choosing an error distribution $\Sigma \leftarrow \Sigma_\alpha$ and adding it to $D_{\mathbf{r}}$, where each $r_i := \alpha((n \cdot d^2)^{1/4} \cdot \sqrt{y_i})$ for $y_1, \ldots, y_{n \cdot d^2}$ sampled from $\Gamma(2, 1)$.

Definition 27. For $i \in [n]$ and a distribution Υ_α over possible error distributions, an algorithm solves the D-CLWE$^i_{q,\Upsilon_\alpha}$ problem if with a non-negligible probability over the choice pairs $(s, \Sigma) \leftarrow U(\Lambda_q^\vee) \times \Upsilon_\alpha$ it has a non-negligible difference in acceptance probability on inputs from $\Pi^i_{q,s,\Sigma}$ and $\Pi^{i-1}_{q,s,\Sigma}$.

This is the average case decision problem relative to R_i; in our worst-case to average-case reduction we will need to randomize the choice of error distribution, which we do by sampling from Υ_α.

Lemma 15. *For any $\alpha > 0$ and $i \in [n]$ there is a randomized polynomial-time reduction from W-D-CLWE$^i_{q,\Sigma_\alpha}$ to D-CLWE$^i_{q,\Upsilon_\alpha}$.*

Proof. Since the definition of Υ_α is a distribution over the family of distributions obtained by sampling from Σ_α and adding an elliptical Gaussian, the proof is the same as Lemma 5.12 of [27], except we replace each instance of mod $\mathfrak{q}_i R^\vee$ with mod R_i and each instance of R_q with Λ_q. □

Remark 8. This choice of Υ_α means that the error covariance matrix in our decision problem is closer to diagonal than that in the corresponding search problem! In fact, if one increased the elliptical error in the decision problem, one could 'flood out' the non-diagonal entries of the covariance matrix, leading to elliptical error which is easier to handle in practice.

Finally, we use a hybrid argument. We must first show that $\Pi^n_{q,s,\Sigma}$ is uniformly random given Σ sampled from Υ_α, but again this follows the same method as the ring case, except we must replace their use of Lemma 1 by [37, Lemma 2.4].

Lemma 16. *Let Υ_α be as above and let $s \in \Lambda_q^\vee$. Then given an oracle $\mathcal{O}$ which solves the D-CLWE$_{q,\Upsilon_\alpha}$ problem there exists an efficient algorithm that solves D-CLWE$^i_{q,\Upsilon_\alpha}$ for some $i \in [n]$ using $\mathcal{O}$.*

Proof. The proof is identical to the ring case, Lemma 5.14 of [27], except that the indexing set $\mathbb{Z}_m^*$ is replaced by $[n]$. □

Denote by CLWE$_{q,\Sigma_\alpha,G}$ the search CLWE problem where $s \in G$ for arbitrary fixed $G \subset \Lambda_q^\vee$. To sum up, we have obtained the main result of this section:

Theorem 7. *Let Λ be the natural order of a cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$, d constant, $q \in poly(n)$ and assume that $\alpha \cdot q \geq \eta_\varepsilon(\Lambda^\vee)$ for a negligible $\varepsilon = \varepsilon(n)$. Then, there is a probabilistic reduction from $CLWE_{q,\Sigma_\alpha,G}$ for any pairwise different $G \subset \Lambda_q^\vee$ to D-$CLWE_{q,\Upsilon_\alpha}$ which runs in time polynomial in n.*

4.3. *Summary and a Remedy for Secret Space*

There are certain technicalities and subtleties in our security proof, which we briefly summarize as follows.

The hardness of Search CLWE in Sect. 4.1 requires a natural order Λ that is maximal. Nonetheless, Lemma 10 (due to Lemmas 6 and 7) is the only stage of the proof that assumes such a natural, maximal order. An improved proof technique may be able to drop this assumption (*e.g.*, to use the natural order). The search to decision reduction in Sect. 4.2 requires a natural order Λ, due to the CRT decomposition of Lemma 12. A better version of CRT may extend the reduction to a maximal order. Fortunately, the orders we take from Theorem 2 are both natural and maximal, thereby meeting these requirements. The requirement of unramified q in Theorem 6 (due to Lemma 6) is minimal: for the algebras of Theorem 2, the only unsuitable primes are the p and q' used in the construction (cf. Sect. 3.3).

Lemma 14 enforces that s lies in a pairwise different set G. It is the only stage of the proof which requires such a set. We emphasize that our reduction takes the search CLWE problem where $s \in G$ for *arbitrary fixed* G to the decision CLWE problem for *arbitrary secret* s. In other words, we claim hardness for the full decision problem, based on hardness of a restricted search problem. Also, our reduction implies that the decision problem is as hard as the search problem for the hardest choice of G. See Appendix D for more details.

Remark 9. The so-called normal form is used de facto in LWE-based cryptography. We note that the normal form reduction is agnostic to the secret space G. More precisely, starting with a secret $s \in G$ gets cancelled in the transformation and replaced by a new secret s' derived from the error distribution (see Lemma 18 in Sect. 5.1). Therefore, the secret space in the normal form of CLWE is the expected space in relation to other LWE normal forms.

Even if our secret space is still exponentially large in n, it may be a concern with security of CLWE if the above reductions were best possible (*e.g.,* decision CLWE is polynomial-time equivalent to restricted search, rather than at least as hard). Fortunately, it is possible to remedy the loss of secret space by using a prime modulus q that totally ramifies in relative extension L/K. The proofs of the following theorems are given in Appendix E.

Theorem 8. *Let $\mathcal{A}$ be a cyclic division algebra over a number field L with center K and natural, maximal order Λ with $|\gamma| = 1$. Let $\alpha = \alpha(n) \in (0, 1)$ and $q = q(n) \geq 2$,*

completely split in K, and the ideals above q in K totally ramify in L, be parameters such that $\alpha \cdot q \geq \omega(\sqrt{\log N})$. Then, there is a polynomial-time quantum reduction from $\mathcal{A}$-DGS$_{\mathcal{I},\xi}$ to search CLWE$_{q,\Sigma_\alpha}$ for any $\xi = r \cdot \sqrt{d}\omega(\sqrt{\log(d \cdot n)})/\alpha q$, where d is constant, $r > \sqrt{2}q \cdot \eta_\varepsilon(\mathcal{I})$ and $\mathcal{I}$ and $q\Lambda$ are coprime.

Note the DGS to search CLWE reduction requires a restriction on the ideal lattice problems that it holds for, but the search to decision part does not depend on any chosen ideal:

Theorem 9. *Let Λ be the natural order of a cyclic division algebra $\mathcal{A} = (L/K, \theta, \gamma)$, d is constant, $q \in poly(n)$ such that the ideals above q in $\mathcal{O}_K$ are maximally ramified in $\mathcal{O}_L$, and assume that $\alpha \cdot q \geq \eta_\varepsilon(\Lambda^\vee)$ for a negligible $\varepsilon = \varepsilon(n)$. Then, there is a probabilistic reduction from CLWE$_{q,\Sigma_\alpha}$ to D-CLWE$_{q,\Upsilon_\alpha}$ which runs in time polynomial in n.*

4.3.1. *Explicit Primes for the Reduction*

Which primes is the reduction valid for? We need $q \in \mathbb{Z}$ such that q splits completely in K, say as $q\mathcal{O}_K = \mathfrak{q}_1 ... \mathfrak{q}_g$, and that these primes are maximally ramified in L, i.e., $\mathfrak{q}_i \mathcal{O}_L = \mathfrak{Q}_i^{[L:K]}$.

To find such primes, we need to review how the algebras used are constructed. We set $K = \mathbb{Q}(\zeta_m)$ and $M = \mathbb{Q}(\zeta_{mq'})$, where $q' = 1 \mod m$ is a prime, and $\gcd(m, q') = 1$. For a degree d extension of K, fix an intermediate field $K \subset L \subset M$ of the correct degree, via the generator of the Galois group of M/K. Recall that we impose $\gcd(d, m) > 1$.

From [45], the ramified primes of $\mathbb{Q}$ in M are the primes dividing mq', and the ramified primes of K in M are the primes dividing q'. Since q' is prime, there is only one prime q dividing it, which is itself. To see that $q = q'$ has the correct ramification, observe the following:

By our choice of q, it is completely split in K. If we label the ramification index e, the inertial degree f, and the number of primes q splits into by g, using the identity $[K : \mathbb{Q}] = e^q_{K/\mathbb{Q}} f^q_{K/Q} g^q_{K/\mathbb{Q}}$, we know that $g^q_{K/\mathbb{Q}} = [K : \mathbb{Q}]$, and $f^q_{K/\mathbb{Q}} = e^q_{K/\mathbb{Q}} = 1$. Moreover, q is ramified in $\mathbb{Q}(\zeta_{mq})$, and q does not divide m. This (with the condition on q) implies that $f^q_{M/\mathbb{Q}} = 1$. Also, $e^q_{M/\mathbb{Q}} = \phi(q) = q - 1 = [M : K]$ and $g^q_{M/\mathbb{Q}} = [K : \mathbb{Q}]$. Multiplicativity of the ramification index and inertial degree then gives $e^q_{L/K} = [L : K]$, $f^q_{L/K} = 1$ and $g^q_{L/K} = 1$, for any intermediate field L.

This means that once an algebra is fixed, there is only one prime that the above reduction is valid for. This might seem like a significant issue; but, to construct an algebra of fixed size, there are infinitely many primes q that can be used to construct M, and thus L. This means that if we know the kind of prime we want to use *before* the algebra is constructed, there are in effect infinitely many primes to choose from. For example, we can consider $K = \mathbb{Q}(\zeta_{128})$, and construct a degree 4 extension of K to generate an algebra of dimension 1024 over $\mathbb{Q}$ using the prime $q = 3457$, and the above reduction holds for those parameters.

5. CLWE in Cryptography

In this section we present a proof-of-concept cryptosystem using CLWE. To demonstrate our comparison against MLWE our scheme will closely resemble the typical 'compact' LWE cryptography schemes over modules, in particular KYBER (see [5]), although it is likely that an adaptation of Regev style encryption from [42] would suit CLWE as well.

5.1. *Making CLWE Suitable for Cryptography: Normal Form*

We implicitly use some standard LWE facts: firstly, we discretize our error distribution e to $\Lambda_q^\vee$; discretizing does not reduce security since an attacker may always discretize the samples themselves. Secondly, we can 'tweak' the problem so that $e, s \in \Lambda_q$. Fortunately, in the case where γ is a unit, $\Lambda^\vee = \bigoplus_i u^i \mathcal{O}_L^\vee$ and so this tweak is precisely multiplying on the right by the tweak factor taking $\mathcal{O}_L^\vee$ to $\mathcal{O}_L$ (see, e.g., [38]). Finally, we require hardness of a 'normal' form for the CLWE distribution, where s is sampled from the same distribution as the noise e.

We require two facts for our proof: firstly, given that q splits completely in K the ring Λ_q is isomorphic to the direct product of n full matrix algebras over $M_{d\times d}(\mathbb{F}_q)$, which can be seen by appealing to the CRT-style decomposition of Lemma 12 and Wedderburn's Theorem as in [33, Propositions 1 and 4]. Secondly, we require that a non-negligible fraction in n of elements of Λ_q are invertible, which follows for fixed, small, d and $q \in \mathrm{poly}(n)$ from this direct product decomposition. Otherwise, our proof follows the outline for that of plain LWE from [4]. Given these two facts, we proceed with showing that the normal form of the CLWE distribution is as hard as the case of taking the secret uniformly at random.

Lemma 17. *For a fixed d and $q \geq (n+1)$, a non-negligible proportion of elements of Λ_q are invertible.*

Proof. Following the decomposition of Lemma 12 and Wedderburn's Theorem, it is sufficient to show that a non-negligible proportion of elements of

$$M_{d\times d}(\mathbb{F}_q) \times \cdots \times M_{d\times d}(\mathbb{F}_q)$$

are invertible, where there are n copies of $M_{d\times d}(\mathbb{F}_q)$. The proportion of invertible elements of $M_{d\times d}(\mathbb{F}_q)$ is precisely

$$\begin{aligned}
&\frac{(q^d-1)(q^d-q)\dots(q^d-q^{d-1})}{q^{d^2}} \\
&= \left(\frac{q^d-1}{q^d}\right)\dots\left(\frac{q^d-q^{d-1}}{q^d}\right) \\
&= \left(1-\frac{1}{q^d}\right)\dots\left(1-\frac{1}{q}\right) \\
&\geq \left(1-\frac{1}{q}\right)^d,
\end{aligned}$$

from which it follows that the total fraction of invertible elements in Λ_q is at least $((1-\frac{1}{q})^d)^n$. By assumption, $q \geq n+1$, and so $(1-\frac{1}{q})^{nd} \geq ((1-\frac{1}{n+1})^n)^d \geq (e^{-1})^d = e^{-d}$, as required. □

Remark 10. This lower bound of e^{-d} means that the normal form reduction will be asymptotic in n but only valid for fixed d. However, as d increases the number of invertible matrices in Λ_q is bounded above by $(1-\frac{1}{q})^{nd}$, and so the reduction would be efficient in d in the case where one enforced a relation on q and d, such as $q \geq nd+1$, or more succinctly $q \geq N$.

Lemma 18. *There is a probabilistic polynomial time reduction from the CLWE problem with uniformly random secret s, possibly over a limited secret space G, and error distribution χ to the CLWE problem with secret $s' \leftarrow \chi$.*

Proof. It is sufficient to show that there is an efficient transformation taking samples with secret s to samples with some new secret s' taken from χ. Sample pairs $(a, b) \leftarrow \Pi_{q,s,\chi}$ until a pair $(a_1, b_1 := a_1 \cdot s + e_1)$ such that a_1 is invertible in Λ_q is obtained. Since a non-negligible fraction of elements of Λ_q are invertible by Lemma 17, this step takes only polynomial time.

Now, given a pair $(a_i, b_i) \leftarrow \Pi_{q,s,\chi}$, we obtain a sample from the CLWE distribution $\Pi_{q,e_1,\chi}$ by outputting $(\overline{a}_i, \overline{b}_i) = (a_i a_1^{-1}, a_i a_1^{-1} b_1 - b_i)$. Since a_1^{-1} is invertible, $\overline{a}_i$ is uniform. Similarly,

$$\begin{aligned} a_i a_1^{-1} b_1 - b_i &= (a_i a_1^{-1}(a_1 \cdot s + e_1)) - a_i \cdot s + e_i \\ &= a_i a_1^{-1} e_1 - e_i, \end{aligned}$$

and so $(\overline{a}_i, \overline{b}_i)$ is a valid CLWE sample with secret e_1 and error distribution χ. Relabeling e_1 as s' completes the proof. □

5.2. *Sample Cryptosystem*

Our scheme is parameterized by an algebra $\mathcal{A} := (L/K, \theta, \gamma)$, where $\mathcal{A}$ is as in Sect. 3.3, an error distribution Σ, and a prime modulus $q \equiv 1 \mod m$ (recall $K = \mathbb{Q}(\zeta_m)$) which is completely split in L. We will denote with bold faced letters the vector form of an element of Λ_q, e.g., if $a = a_0 + ua_1 + ... + u^{d-1}a_{d-1}$ then $\mathbf{a} = (a_0, a_1, \ldots, a_{d-1})$. We note that $\mathcal{O}_L/q\mathcal{O}_L$ has a polynomial representation of dimension $n \cdot d$, and so we encode our message $\in \{0, 1\}^{n \cdot d^2}$ as an entry of Λ_q as a vector $\mathbf{m}$ of d $\{0, 1\}$ polynomials. The scheme proceeds as follows:

- Alice generates a CLWE sample $(a, b := a \cdot s + e)$, where $a \in \Lambda_q$ is uniformly random and $s, e \leftarrow \Sigma$, and outputs public key $\mathbf{a}, \mathbf{b}$.
- To encrypt $\mathbf{m} \in \{0, 1\}^{n \cdot d^2}$, Bob samples $t, e_1, e_2 \leftarrow \Sigma$ and outputs $\mathbf{u} := \phi(a)^T \mathbf{t} + \mathbf{e}_1$, $\mathbf{v} := \phi(b)^T \mathbf{t} + \mathbf{e}_2 + \lceil \frac{q}{2} \rfloor \cdot \mathbf{m}$.
- To decrypt, Alice computes $\mathbf{c} = \mathbf{v} - \phi(s)^T \mathbf{u}$ and recovers each coordinate of $\mathbf{m}$ by rounding the corresponding entry of $\mathbf{c}$ to 0 or $\lceil \frac{q}{2} \rfloor$ and outputting 0 or 1 respectively.

Remark 11. There are two benefits of instantiating this scheme in the cyclic algebra setting rather than over modules as in [5], both following from the matrix embedding ϕ. Firstly, in the module setting Alice must publish a matrix A rather than the vector $\mathbf{a}$ in her key, since $\phi(a)$ lets us generate a matrix; this saves a factor of d in the size of the public key. Secondly, by extending $\mathbf{b}$ to $\phi(b)$ we are able to increase the dimension of $\mathbf{v}$, and correspondingly increase the size of the message by a factor of d.

Example 3. Recall our explicit algebras from Sect. 3.3. Without considering streamlined implementation for specific NIST submissions, we will pick toy comparison parameters for equivalent module-based systems and ring-based schemes, e.g., KYBER and NewHope. For the module case, consider a module of dimension 4 over a ring L of dimension 256, with 2-power cyclotomic base field $[K : \mathbb{Q}] = 64$. Our public key $(\mathbf{a}, \mathbf{b})$ requires storing only 8 elements of $R_q = \mathcal{O}_L/q \cdot \mathcal{O}_L$ rather than 20 in the form $(A, \mathbf{b})$. Our message consists of 1024 bits, corresponding to the total dimension of the algebra rather than the module versions 256 which corresponds to the field dimension; if the private key size is 256, our CLWE scheme allows a rate-1/4 binary error correction code, while KYBER does not. Our ciphertext sizes are the same. As far as the modulus q is concerned, we find $q = 3329$ splits completely in a quartic cyclic extension L of K, which matches with the modulus q used in KYBER;[7] meanwhile, $q = 3457$ splits completely in K but ramifies totally in another relative extension of K. Overall this represents a noteworthy gain in key and message size without loss in efficiency. For the ring case, consider an instantiation of NewHope in dimension 1024. Both public keys are in the form (a, s) and so require equivalent levels of storage (8 elements of a field of dimension 256 or 2 in dimension 1024), and the same phenomenon is true of ciphertext sizes and message length. However, a larger modulus $q = 12289$ is used in NewHope. Hence, we hope to gain in security without losing much efficiency. A limitation of our current method is that we cannot achieve rank $d = 3$, similar to the RLWE limitation over power-of-2 rings.

Before considering security and correctness we need a somewhat technical lemma allowing the use of the matrix transpose operation. Essentially, it states that if the CLWE problem is hard in an algebra $\mathcal{A}$, then for $a, s, e \in \Lambda_q$, the equation $\phi(a)^T\mathbf{s} + \mathbf{e}$ is a valid CLWE instance in some other algebra $\mathcal{A}'$ for which the CLWE problem is still hard.

Lemma 19. *Let $\mathcal{A} = (L/K, \theta, \gamma)$, where γ is a unit, be a cyclic division algebra with matrix embedding $\phi(a)$ and natural order Λ. Then there exists another cyclic algebra $\mathcal{A}' = (L/K, \theta, \gamma^{-1})$ with matrix embedding $\phi'(a')$ and natural order Λ' such that for $a \in \mathcal{A}$ there exists $a' \in \Lambda'$ satisfying $\phi(a)^T = \phi'(a')$. Moreover, $\mathcal{A}'$ still satisfies the division algebra condition, and Λ'_q and Λ_q canonically isomorphic as additive groups.*

Proof. The fact that $\mathcal{A}'$ is still a division algebra follows from the non-norm property on γ and the fact that $N_{L/K}(L^\times)$ is a multiplicative group. Λ'_q and Λ_q are additive isomor-

[7] The initial version of KYBER uses $q = 7681$, but it has been reduced to 3329 later which does not split completely in $L = \mathbb{Q}(\zeta_{512})$. It is noteworthy that, with a similar technique, further reduction of q in CLWE may also be possible.

phic because both algebras share the same underlying fields and γ, γ^{-1} are both units of $\mathcal{O}_L$. Since the first row of $\phi(a)$ is precisely $(x_0, \gamma\theta(x_{d-1}), \gamma\theta^2(x_{d-2}), \ldots, \gamma\theta^{d-1}(x_1))$, by setting $a' = x_0 + u\gamma\theta(x_{d-1}) + \cdots + u^{d-1}\gamma\theta^{d-1}(x_1)$ and observing that θ^d is the identity it is easy to check that $\phi(a)^T = \phi'(a')$. □

The proofs of correctness and security are similar in spirit to those of other compact LWE schemes such as, e.g., NewHope [3] or KYBER [5]. We proceed with a somewhat informal security argument.

Lemma 20. *The defined scheme is IND-CPA secure under the assumption that the decision $CLWE_{q,\Upsilon}$ problem is hard.*

Proof. The goal of an IND-CPA adversary is to distinguish, with non-negligible advantage, between encryptions of two plaintexts m_1, m_2. The challenger chooses $i \in \{0, 1\}$ uniformly at random and encrypts m_i as $\mathbf{u}, \mathbf{v}$. By the assumption that the decision CLWE problem is hard, the adversary cannot distinguish between the case where $b = as + e$ and the case where it is replaced by a uniform random b', so we replace b in the public key given to the adversary by b' and also use b' to compute the challenge ciphertext $\mathbf{v}'$. Setting $\mathbf{v}'' := \mathbf{v}' - \lceil \frac{q}{2} \rfloor \cdot \mathbf{m}_i$, it follows by Lemma 19 that $\mathbf{u}, \mathbf{v}''$ represent two samples from a valid CLWE distribution with secret $\mathbf{t}$, and so the adversary cannot distinguish them from uniform with non-negligible advantage. Hence, the challenger cannot distinguish $\mathbf{v}'$ and hence $\mathbf{v}$ from uniform with non-negligible advantage and so cannot guess i with non-negligible advantage. □

Finally, we demonstrate conditions on the error term for the scheme to be correct.

Lemma 21. *The defined scheme is correct as long as the ℓ_∞ norm of $\boldsymbol{e}' = (\phi(e)^T \boldsymbol{t} + \boldsymbol{e}_2 - \phi(s)^T \boldsymbol{e}_1)$ is less than $\lceil \frac{q}{4} \rfloor$, where the ℓ_∞ norm is over the vector of all polynomial coefficients of each u^i entry of $\boldsymbol{e}'$ of dimension $n \cdot d^2$.*

Proof. To decrypt, Alice computes $\mathbf{v} - \phi(s)^T \mathbf{u}$ and computes $\mathbf{m}$ by rounding. Since $\phi(\cdot)$ is a homomorphism, we have

$$\begin{aligned}
\mathbf{v} - \phi(s)^T \mathbf{u} &= \phi(b)^T \mathbf{t} + \mathbf{e}_2 + \left\lceil \frac{q}{2} \right\rfloor \cdot \mathbf{m} - \phi(s)^T (\phi(a)^T \mathbf{t} + \mathbf{e}_1) \\
&= \phi(e)^T \mathbf{t} + \mathbf{e}_2 - \phi(s)^T \mathbf{e}_1 + \left\lceil \frac{q}{2} \right\rfloor \cdot \mathbf{m} \\
&= \mathbf{e}' + \left\lceil \frac{q}{2} \right\rfloor \cdot \mathbf{m}.
\end{aligned}$$

from which the result follows immediately. □

We note that the error term $\mathbf{e}'$ will be unsurprising to those familiar with LWE-based cryptography. Although we do not provide concrete correctness estimations, the error parameters for our decision reduction are equivalent to those of MLWE up to some small covariance terms. We do not expect this covariance to greatly affect the distribution of the

error and thus for equivalent parameter choices we expect a similarly small probability of decryption failure.

5.3. *Operational Complexity in Cyclic Algebras*

In the previous subsection we showed that the CLWE problem can be used to construct a standard LWE-based cryptosystem. Assuming that parameters across all variants of the LWE assumption are roughly equivalent, the CLWE problem supports key and message sizes as advantageous as those of the RLWE problem, and better than those of the module case. Along with storage considerations, another important facet of the ambient space in LWE cryptography is the efficiency of operations. Here, we will consider the asymptotic complexity of multiplication in a cyclic algebra in order to compare it to the ring and module variants. Since in practice we consider operations modulo some prime q, addition in rings, modules, and cyclic algebras can be considered as addition in vector spaces over $\mathbb{Z}_q$, which has complexity dominated by that of multiplication.

Consequently, we only concern ourselves with a comparison of the cost of computing the multiplication operation $A\mathbf{s}$ in the three cases. In order to keep our comparison consistent, we let N denote the total dimension of the underlying LWE instance. In the ring case, N denotes the ring dimension; in the module case, $N = nd$, where n denotes the ring dimension and d the module rank; in the cyclic algebra case $N = nd^2$, where the ring dimension is nd and the algebra has 'module' rank d. However, since it will be important later we remark here that the cyclotomic part of the ring will be of dimension n rather than nd. The three cases can be considered as follows:

- In the ring case, the operation $A\mathbf{s}$ over $\mathbb{Z}_q$ is a representation of the ring operation $a \cdot s$ in $R_q \cong \mathbb{Z}_q[X]/(X^N + 1)$. Using the CRT decomposition in dimension N of [28], this operation is decomposed into coordinatewise multiplication in a vector of dimension N over $\mathbb{Z}_q$, following which the decomposition is reversed to recover $a \cdot s$. The complexity of this technique is dominated by that of the CRT decomposition, which takes time $O(N \log N)$, although the coordinatewise multiplication also requires time $O(N)$.
- In the module case, A is a $d \times d$ matrix over R_q. In this case, one can compute $A\mathbf{s}$ by applying the CRT in dimension n coordinatewise on A and $\mathbf{s}$. This requires $d^2 + d$ applications of the CRT, for a total asymptotic complexity of $O(d^2 n \log n) = O(Nd \log(N/d))$. Again, this hides a coordinatewise multiplication step which takes time $O(Nd)$ in this setting.
- In the cyclic algebra case, A is a matrix in the shape $\phi(a)$, where $\phi(a)$ is the left regular representation of $a \in \Lambda_q$. We estimate the complexity of the operation $\phi(a) \cdot \mathbf{s}$ in Appendix F. Explicitly, our algorithm has complexity $O(N \log(N/d^2)) + \tilde{O}(Nd^{\omega-2})$ in the case where q splits completely in L, with $\omega \in [2, 2.373]$ denoting the exponent of matrix multiplication. The latter term corresponds to the cost of multiplication in our analog of the finite fields used in the CRT method for RLWE.

We see that, in the case of completely splitting q, cyclic algebras compare favorably with modules for multiplication in the same dimension N and when d grows to infinity, depending on the exact relationship between $\log d^2$ and $d^{\omega-2}$. Recall that for our reduction to hold, we require d to be constant, in which case all three complexities discussed

above are the same (since the constant d will be hidden by the constant appearing in the $O(\cdot)$). Moreover, we currently do not know how to construct CLWE instances for arbitrary field degree and module rank, e.g., $n = 256$ and $d = 3$ like in Kyber.

6. Conclusions and Future Work

The primary goal of this work is the introduction of the Learning with Errors problem over Cyclic Algebras, CLWE, adding to the family of available LWE assumptions for use in cryptography. To this end, the central pillars of an LWE problem are provided for the cyclic algebra case. First, in order to provide a foundation for the construction the notion of lattices derived from two-sided ideals of the natural order of a cyclic algebra are applied in cryptography for the first time. Then, in Sect. 3, the CLWE problem is formally introduced, following which explicit algebras are provided with dimensions and structure appropriate for cryptographic use. Then, in Sect. 4, the usual LWE security reductions are established in the CLWE case, namely, samples from the CLWE distribution appear pseudorandom to an onlooker with no knowledge of the secret s. Finally, in Sect. 5, the necessary steps are taken to mold the CLWE problem into a practical format for cryptography. Normal form reduction is shown and a sample cryptosystem in this form is provided. Additionally, the complexity of operations in CLWE cryptography is compared to that of RLWE and MLWE-based schemes.

Cyclic algebras exhibit substantial novel structures within lattice-based cryptography, and discovering use cases for these previously unseen features represents an exciting area of future research. We outline a few directions of future research in the following.

From a theoretical standpoint, the most pressing question to be solved about CLWE is whether or not the search and decision problem are polynomial time equivalent, or instead if the hardness of the decision variant can be based directly on hard lattice problems via some other technique. In this work, the effectiveness of our technique to show the hardness of the decision problem depended on the modulus q: the case of completely split q resulted a loss of secret space; while the case of ramified q remedied this issue, we have not managed to come up with efficient multiplication.

Another method of establishing the hardness of decision RLWE that is not shown for CLWE in this work is a direct to decision reduction, which more generally represents a security proof for the decision problem that holds for wider classes of cyclic division algebras than those of Sect. 4.2. The direct to decision reduction of [40] is the only security reduction for RLWE which establishes the hardness of the decision problem without enforcing that K is a cyclotomic field within which q splits completely, as in the search-decision reduction of [27] and the presented analog for CLWE. Dropping this restriction, and hence widening the possible choices of cyclic algebras supporting the hardness of the decision problem, would provide larger design space for CLWE-based cryptography.

As for another direction of future work, we view a drawback of our work to be that we are restricted to certain instances of cyclic algebras. Although in practice most cryptography would use a fixed choice of algebra, this is a function of our methods and may be possible to remove. Additionally, showing the aforementioned direct-to-decision reduction may generalize the choice of algebras.

Finally, this work is focused on the theoretical construction of a non-commutative Ring-LWE assumption, and we leave practical analysis and implementation of cryptography based on CLWE as further research.

Acknowledgements

The authors would like to thank Jyrki Lahtonen, Damien Stehle and Martin Albrecht for helpful discussions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

A. Impossible Algebras

We show that certain algebras that would otherwise be what we are looking for do not exist under our restrictions. As discussed above we would like to begin with a base field that is cyclotomic, $K = \mathbb{Q}(\zeta_m)$ for integer m, and proceed to fix some low degree cyclic Galois extension L/K and non-norm element $\gamma \in \mathcal{O}_K$ with $|\gamma| = 1$, e.g., γ is a root of unity. Given these restrictions and the shape of lattice cryptography, the most natural fields to look for are low degree extensions of two-power cyclotomics, e.g., $m = 2^k$. Unfortunately, we are able to prove the non-existence of a large class of such extensions.

Theorem 10. *Let $K = \mathbb{Q}(\zeta_m)$ for some positive integer m and let $p \geq 2$ be some integer which is coprime with m. Then, for any Galois extension L/K of degree p each $\zeta_m, \zeta_m^2, \ldots, \zeta_m^{m-1}$ lies in $N_{L/K}(K^\times)$.*

Proof. Since L/K is a Galois extension of degree p, the relative norm map $N_{L/K}(\cdot)$ induces the map $x \to x^p$ on elements $x \in K^\times$. Let $1 \leq i \leq m-1$ be an integer; we will prove the theorem by finding $1 \leq j \leq m-1$ such that $N_{L/K}(\zeta_m^j) = \zeta_m^i$. Since ζ_m and its powers lie in K, the relative norm map takes ζ_m^j to ζ_m^{jp} and we are left to solve the congruence $jp \equiv i \mod m$. By assumption, g.c.d.$(m, p) = 1$ and so p is invertible modulo m. Denoting this inverse p^{-1} and letting $j = p^{-1}i \mod m$ it is easy to see that $jp \equiv ip^{-1}p \equiv i \mod m$. The theorem statement follows immediately. □

This theorem precludes the existence of a very large class of cyclic division algebras with cyclotomic base field. In particular, if the degree of $[L : K]$ is coprime with m then we cannot have our restrictions that $|\gamma| = 1$, is integral, and that K is cyclotomic. We draw attention to the specific classes whose non-existence we are interested in: in an

ideal world we might instantiate CLWE with $K = \mathbb{Q}(\zeta_{2^k})$ and $[L : K] = d$ for arbitrary small integer d corresponding to the module rank, which in practice is likely to be at most say 5. However, as a result of Theorem 10 we know that d cannot be coprime with 2^k and must be even in order to permit a suitable γ, from which it follows that we cannot have $d = 3, 5$.

B. Proofs of Theorem 3 and Theorem 4

Before proving Theorem 3 we need some additional concepts and a Lemma. Given a K-central division algebra $\mathcal{A}$ and some $\mathcal{O}_K$ order Λ in it, then the $\mathcal{O}_K$*-discriminant* of Λ, $d(\Lambda/\mathcal{O}_K)$, is a certain ideal in $\mathcal{O}_K$ [43, p.126]. While $\mathcal{A}$ has many maximal orders they all share the same discriminant, which is called the discriminant of the algebra $d_{\mathcal{A}}$. Now the key fact about discriminants we need is that an order Λ is maximal if and only it's discriminant equals that of $d_{\mathcal{A}}$.

We will now use the notation of Sect. 3.3. According to [21] the field M and therefore also its subfield L are subfields of $\mathbb{Q}(\zeta_m, \zeta_{q'})$, where $m = p^a$, and $q' \neq p$ is some large prime. Let $n = \varphi(m) = p^{a-1}(p-1)$. Furthermore it is known that q' splits completely in the field $K = \mathbb{Q}(\zeta_m)$. Let us now denote with

$$q'\mathcal{O}_K = \mathfrak{q}'_1 \dots \mathfrak{q}'_n,$$

the prime ideal decomposition of q' in K. We then have the following result.

Lemma 22. *Let $(L/K, \theta, \zeta_m)$ be an index d division algebra of Theorem 2 and let Λ be the corresponding natural order. Then we have that*

$$d(\Lambda/\mathcal{O}_K) = (\mathfrak{q}'_1, \dots \mathfrak{q}'_n)^{d(d-1)}. \tag{2}$$

Proof. According to [44, Lemma 5.4] we have that

$$d(\Lambda/\mathcal{O}_K) = d(L/K)^d \zeta_m^{d(d-1)} = d(L/K)^d,$$

where $d(L/K)$ is the relative number field discriminant of the extension L/K. In order to find the discriminant of the natural order, it is now enough to find $d(L/K)$. By the basic theory of cyclotomic fields we know that $\mathbb{Q}(\zeta_m, \zeta_{q'}) = \mathbb{Q}(\zeta_{mq'})$. We also know that the only ramified primes in the extension $\mathbb{Q}(\zeta_{mq'})/\mathbb{Q}$ are p and q' and their ramification indices are $e_1 = n$ and $e_2 = q' - 1$, respectively. Furthermore ramification index of p in the extension $\mathbb{Q}(\zeta_m)/\mathbb{Q}$ is e_1. As ramification indices are multiplicative in towers of extensions we can deduce that the only primes that are possibly ramified in the extension $\mathbb{Q}(\zeta_{mq'})/\mathbb{Q}(\zeta_m)$ are those that lie above q' in the ring $\mathcal{O}_K$. As q' is not ramified in $\mathbb{Q}(\zeta_m)$, we get again by the multiplicativity of the ramification indices that all the primes $\mathfrak{q}'_i$ are totally ramified in the extension $\mathbb{Q}(\zeta_{mq'})/\mathbb{Q}(\zeta_m)$. Therefore, they are also totally ramified in the extension $L/\mathbb{Q}(\zeta_m)$. Because q' does not divide d the prime ideals $\mathfrak{q}'_i$ are tamely

ramified. Dedekind's discriminant theorem now implies that

$$d(L/K) = (\mathfrak{q}'_1 \dots \mathfrak{q}'_n)^{(d-1)}.$$

□

Now we are ready to prove the natural order in Theorem 3 is actually maximal.

Proof. The proof is based on the result in [43] that states that an order is maximal if and only if it has the same discriminant as the discriminant of the algebra. According to Lemma 22 we have that

$$d(\Lambda/\mathcal{O}_K) = d(L/K)^d = (\mathfrak{q}'_1 \dots \mathfrak{q}'_n)^{d(d-1)}. \tag{3}$$

According to [43] the discriminant of the maximal order will always divide the discriminant of the natural order. Hence, we know that the only prime ideals that can possibly divide the discriminant of the maximal order are $\mathfrak{q}'_i$. Let us now assume that $\mathfrak{Q}_i$ is prime ideal above $\mathfrak{q}'_i$ in M. By abusing notation we will denote with $M_{\mathfrak{q}'_i}$ the $\mathfrak{Q}_i$-adic completion of M and in the same way the respective completion $L_{\mathfrak{q}'_i}$.

Following the proof of [21, Theorem 4] we can see that the authors actually prove that ζ_m is a non-norm element in the extension $M_{\mathfrak{q}'_i}/K_{\mathfrak{q}'_i}$ for each prime ideal $\mathfrak{q}'_i$. Using the same proof as in Theorem 2 we can now see that ζ_m is a non-norm element in the extensions $L_{\mathfrak{q}'_i}/K_{\mathfrak{q}'_i}$, for all i. According to [43, Theorem 30.8] $A \otimes_K K_{\mathfrak{q}'_i} \cong (L_{\mathfrak{q}'_i}/K_{\mathfrak{q}'_i}, \theta', \zeta_m)$, where θ' naturally extends θ. As ζ_m is a non-norm element, $(L_{\mathfrak{q}'_i}/K_{\mathfrak{q}'_i}, \theta', \zeta_m)$ is an index d division algebra. By definition of the local index we can see that the local indices $m_{\mathfrak{q}'_i}$ are d for all $\mathfrak{q}'_i$. We now know that $\mathfrak{q}'_i$ are the only possible primes dividing the discriminant and that their local indices are d. According to [43, Theorem 32.1] the discriminant of the algebra $\mathcal{A}$ is

$$d_{\mathcal{A}} = \prod_{i=1}^{n} \mathfrak{q}'^{\,(m_{\mathfrak{q}'_i}-1)\frac{d^2}{m_{\mathfrak{q}'_i}}}_i = \prod_{i=1}^{n} \mathfrak{q}'^{\,(d-1)d}_i,$$

completing the proof. □

The proof of Theorem 4 is similar.

Proof. We have $K = \mathbb{Q}(\zeta_n)$ for $n = p^r$ where p is prime, and L/K a degree d extension. Thus, we have $K \subset L \subset \mathbb{Q}(\zeta_{nm})$ for some integer m, by the Kronecker-Weber theorem. In our context, we may take $gcd(n, m) = 1$. The prime ideals of $\mathcal{O}_{KF}$ which ramify in $\mathcal{O}_{LF}$ lie above the same integer primes as the prime ideals of $\mathcal{O}_K$ which lie above $\mathcal{O}_L$, because of the disjointness of L and F. Denote this set of primes by $S = \{p_1, \dots, p_l\}$. Write $p_i\mathcal{O}_{KF} = \prod_j \mathfrak{p}_{ij}$; the ramification index of p_i in KF is 1. Moreover, $\mathfrak{p}_{ij}$ is totally ramified in $\mathbb{Q}(\zeta_{nmq^t})$, and if ramified in LF, is totally ramified in LF by multiplicativity of the ramification data. Since L/K induces a CDA with maximal natural order, and $[LF : KF] = [L : K]$, we know that $p_i \nmid d$, and so the

$\mathfrak{p}_{ij}$ are tamely ramified. This means we can apply Dedekind's discriminant theorem and obtain $d(\Lambda'/\mathcal{O}_{KF}) = d(LF/KF)^d = \prod_{i,j} \mathfrak{p}_{ij}^{d(d-1)}$.

It remains to see that $\mathcal{A}'$ is a division algebra, and what $d_{\mathcal{A}'}$ is. As in the subfield case, we consider algebras arising from completions of number fields at certain prime ideals. Let $p_i\mathcal{O}_K = \prod_j \mathfrak{q}_{ij}$. We know that by construction $A_{\mathfrak{q}_{ij}} = (L_{\mathfrak{q}_i j}/K_{\mathfrak{q}_i j}, \theta^*, \zeta_n)$ is an index d CDA, where $\mathcal{A}_{\mathfrak{q}_{ij}}$ denotes the completion of $\mathcal{A}$ at $\mathfrak{q}_{ij}$ and θ^* extends θ. We are interested in the index of algebras $A'_{\mathfrak{p}_{ij}} = (LF_{\mathfrak{p}_{ij}}/KF_{\mathfrak{p}_{ij}}, \theta'', \zeta_n)$. These can be presented in form $A_{\mathfrak{q}_{ij}} \otimes_{K_{\mathfrak{q}_{ij}}} KF_{\mathfrak{p}_{ij}} \cong A'_{\mathfrak{p}_i j}$. It is a consequence of [43, Theorem 31.9] that $A'_{\mathfrak{p}_i j}$ has local index d if and only if $gcd([KF_{\mathfrak{p}_{ij}} : K_{\mathfrak{q}_{ij}}], d) = 1$. As KF/K is a Galois extension we know that $[KF_{\mathfrak{p}_{ij}} : K_{\mathfrak{q}_i j}]$ divides $[KF : K] = [F : \mathbb{Q}]$. Therefore, it follows that since $gcd(d, [F : \mathbb{Q}]) = 1$, also $gcd([KF_{\mathfrak{p}_{ij}} : K_{\mathfrak{q}_{ij}}], d) = 1$. We can conclude that $m_{\mathfrak{p}_{ij}} = d$ for all $\mathfrak{p}_{ij}$. It follows that A' is a division algebra and that $d_{A'} = \prod_{i=1} \mathfrak{p}_{ij}^{(m_{\mathfrak{p}_{ij}}-1)\frac{d^2}{m_{\mathfrak{p}_{ij}}}} = \prod_{i=1}^{n} \mathfrak{p}_{ij}^{(d-1)d} = d(\Lambda'/\mathcal{O}_{KF})$, as required. □

C. Extensions where q Splits Completely in L

We would like q to be of roughly appropriate cryptographic size (say between 3000 and 15000 as a soft estimate, once again presuming parameters similar to those of NewHope or KYBER). Having q split completely in L is not as straightforward as in K because L is not a cyclotomic field, so we return to our examination of the proof of Theorem 1. Recall that in this proof the extension field L is a subfield of $K(\zeta_{mq'})$ for some prime integer q' satisfying $q' = 1 \mod m$ and, for $m = p^a$, p^{a+1} does not divide $q' - 1$. That is, a is the highest power of p that divides $q' - 1$. We have several methods to ensure that q splits completely in L, of which we start with the most naive.

Naive Method For our general method we rely on the following fact: If $\mathfrak{q}_i$ is an ideal of $\mathcal{O}_K$ which splits completely in an extension M/K then it splits completely in any intermediate field $M/L/K$. As it is conceptually simpler to apply this idea to the integer q than to the $\mathcal{O}_K$-ideals $\mathfrak{q}_i$, we use a simpler statement, that if $\langle q \rangle$ splits completely in some M containing L then it splits completely in L. This gives us an easy way to find some q that splits completely by examining a cyclotomic field that contains L: let $K = \mathbb{Q}(\zeta_m)$ and let $M = K(\zeta_{q'})$. Then since $q' = 1 \mod m$ it follows that $M = \mathbb{Q}(\zeta_{mq'})$. Thus, q splits completely in M if and only if $q = 1 \mod mq'$ and consequentially splits completely in our extension L if $q = 1 \mod mq'$. Since there are infinitely many primes equal to $1 \mod mq'$ this recipe always provides a prime q that splits completely in L. The upside of this method is that it is both very general and simple, since all candidate fields L we construct are contained in a larger cyclotomic field. Theoretically, this method can be extended to any abelian extension of $\mathbb{Q}$ using the partial converse of the Kronecker–Weber Theorem. However, using the Kronecker–Weber Theorem constructively is not as straightforward as picking q' as in the proof of Theorem 1, so this extension to general abelian L is slightly contrived.

The downside to this method is that it seems that often this will result in unrealistically large q. Since $q' = 1 \mod m$ and not $1 \mod p^{a+1}$, q' must be chosen carefully

and there are not many 'small' primes satisfying these conditions. For example, in our quadratic extension case with $m = 512$ the smallest prime that is $1 \mod m$ but not $1 \mod 2m$ is $q' = 7681$. The smallest q which is $1 \mod (512 \cdot 7681)$ has to be bigger than $512 \cdot 7681 = 3932672$, which is inappropriately large for lattice cryptography. Of course, one could be lucky here and have much smaller q for different choices of L and K, but in general we regard this as a theoretical result rather than a practical method. Even for smaller 2-power cases such as $m = 128$ one must set $q' = 641$, which leads to a smallest valid prime of $q = 820481$.

Remarkably, this is much less bad in the cubic case; $K = \mathbb{Q}(\zeta_{81})$ gives $q' = 163$ as a suitable prime and $q = 26407$ still splits completely. This is perhaps slightly too large, but certainly not so much so that it is completely impractical. Nonetheless, we move on to a better method for quadratic cases.

Quadratic case In the case where L/K ($K = \mathbb{Q}(\zeta_{512})$) is a quadratic extension we are able to choose substantially smaller q by examining the unique quadratic subfields of $E' := \mathbb{Q}(\zeta_{q'})$. We rewrite M as the compositum of E' and K, and observe that since our chosen L contains K our method of choosing L as a subfield of M allows us to write $L = EK$ for a subfield E of E'. In the case where L is a degree two extension of K we know that E is a quadratic field, and since E' is a prime cyclotomic field we have an explicit description for its unique quadratic subfield E; namely that $E = \mathbb{Q}(\sqrt{q'})$ if $q' = 1 \mod 4$ and $E = \mathbb{Q}(\sqrt{-q'})$ is $q' = 3 \mod 4$. It is a standard fact that the discriminant d_E of E is q' if $q' = 1 \mod 4$ and $-q'$ otherwise. Finally, we know that a prime q splits completely in E if and only if the congruence $d_E = x^2 \mod q$ has a solution, e.g., if d_E is a square mod q. Plugging in the prime numbers $q = 12289$ and $q' = 7681$ that are common in cryptography we see that $q' = 1 \mod 4$ and that $7681 = 3788^2 \mod 12289$, so that $q = 12289$ splits completely in E, K, and thus L, as required. Since this prime is explicitly the prime used in NewHope for all parameter sets we view this method as a substantial improvement on the previous technique.

Quartic fields Again, we use the method of describing L as a compositum MK/K. Now, M will be a quartic subfield of the field $\mathbb{Q}(\zeta_{q'})$ and one can establish the linearly disjoint nature of M and K required to express L as this compositum by, e.g., examining their discriminants: since K is a power-of-two cyclotomic field the only prime appearing in its discriminant is 2, and since M is a subfield of $\mathbb{Q}(\zeta_{q'})$ the only prime in its discriminant is q'. Since they have coprime discriminants they are linearly disjoint, and since ramified primes are factors of the discriminant we have a relatively easy way to discount q being ramified ($q \neq 2, q'$), so the remaining case to concern ourselves with is q being inert.

Since the discriminants are coprime we have a method for explicitly describing the integral basis of $L = MK$; the integral basis for K is clear, and an integral basis for M in fixed dimension can be computed relatively easily since it has degree 4. Then, the product of their integral bases is an integral basis for L. Now one only needs to check whether q splits completely in M, since splitting in K is well understood. We are unable to provide a general method for finding such q, but an easy computation reveals that for $q = 10753$ and $K = \mathbb{Q}(\zeta_{256})$ there is a quartic field M such that q splits completely in M and K and hence L. Since we have a relatively small range in which we wish to place q and M has low degree we do not consider the cost of this search as a large

drawback since it can be done efficiently on computational software such as SAGE or PARI.

Remark 12. In fact, this quartic method can be applied to other instances where we do not have an explicit description of the subfields of $K(\zeta_{q'})$ which have degree d over K: define the families of q which split completely in K, then check whether those q split completely in L using computational software. Since $q = 1 \mod m$ and m is relatively large, there will not be many q to check of appropriate size for lattice cryptography, and so we conclude that this method is sufficient for fixed choices of fields L, K for which a satisfactory q exists.

Compositum Fields Since a prime q is completely split in a compositum field K_1K_2 if and only if it is completely split in both K_1 and K_2, it is ready to extend the above method to compositum fields.

For the case of Fig. 3a, suppose we have found primes q completely split in K' and L' using the above method. Then we choose q that is also completely split in F, which ensures it is completely split in compositum field $K = K'F$, hence in $L = L'F$.

For the case of Fig. 3b, we choose q that is also completely split in K, which ensures it is completely split in compositum field $L = KL'$.

D. Restricting the Secret Space

In Lemma 14 we need to use a fact that is implicit in the search-decision reduction of [27]: for uniformly random $v \in \mathcal{R}_i$ and an incorrect guess g of the secret s modulo $\mathcal{R}_i$, the distribution of $v(g-s)$ is uniformly random. In the ring and module cases, the secret space is decomposed into a direct product of finite fields, so it is clear that $v(g-s)$ is uniformly random in each finite field for $g \neq s$.

In our case, an appeal to Wedderburn's theorem demonstrates that, since for our parameter choices each $\mathcal{R}_i$ is a central simple algebra over ${\mathcal{O}_K}^\vee/\mathfrak{q}_i{\mathcal{O}_K}^\vee \cong \mathbb{F}_q$, each $\mathcal{R}_i$ is isomorphic to the full matrix ring $M_{d\times d}(\mathbb{F}_q)$, for which it is not true in general that $v(g-s)$ is uniformly random for $g \neq s$; in fact, it is uniformly random if and only if $g - s$ is invertible. Thus, we restrict our secret s so that $s \mod \mathcal{R}_i$ lies in a set G_i with the property that $g \neq h \in G_i$ implies $g - h$ is an invertible matrix. Applying this restriction for each i places $s \in G$ for a set $G = G_1 \times \cdots \times G_n$ of size $|G| = \prod_i |G_i|$. Now, an incorrect guess $g \in G_i$ of $s \mod \mathcal{R}_i$ results in a distribution of $v(g-s)$ which is uniformly random mod $\mathcal{R}_i$. We will call such a set G a pairwise difference set.

We also need to guarantee that there exist sufficiently large choices of G. A simple method for constructing a valid G_i is by fixing some arbitrary embedding β of $\mathbb{F}_{q^d}$ into $M_{n\times n}(\mathbb{F}_q)$ and letting G_i equal the image of this embedding, such that $|G_i| = q^d$ and $|G| = q^{nd}$. Indeed, a G_i constructed in this way is maximal because any set of matrices in $M_{d\times d}(\mathbb{F}_q)$ of size at least $q^d + 1$ contains two matrices with the same first row, whose difference is therefore uninvertible.

There are a number of choices of embedding β, and thus set G_i, equal to the number of irreducible polynomials of degree d in $\mathbb{F}_q[x]$, which can be calculated by the Necklace polynomial and in general will vastly exceed q. We make clear that our

reduction will take the decision CLWE problem for *arbitrary secret* s to the search CLWE problem where $s \in G$ for *arbitrary fixed* G, which we denote by $\mathrm{CLWE}_{q,\Sigma_\alpha,G}$. Thus, our reduction states that the decision problem is as hard as the search problem for the hardest choice of G, precluding obvious attacks on the unique case where $G = \mathcal{O}_{Lq}{}^{\vee}$ and the CLWE problem with $s \in G$ corresponds to d parallel copies in L of the RLWE problem.[8] For a general set G, $s \in G$ will not provide parallelization since they need not have the property of L that they are entirely contained in one u coordinate of $\mathcal{A}$. Additionally, even though elements of G constructed this way co-commute, they do not lie in the center of Λ and the multiplication $a \cdot s$ in the CLWE instance will not be a commutative operation.

Of course, fixing a G of size q^{nd} restricts the size of the secret space by a factor of $\frac{q^{nd}}{q^{nd^2}}$, a substantial loss in size even for fixed, small d. For concrete parameter settings, this may result in a much easier problem, but asymptotically it is still exponential in n and thus establishes a suitable hardness property for decision CLWE. Of course, attacks based on exhaustive search are unlikely to represent the best attacks on the CLWE problem, so this may or may not substantially aid an attacker in practice.

In fact, there is no a priori reason why G_i should be a field, or even closed under multiplication. For example, fixing a pair of invertible matrices M_1, M_2 and replacing G_i with $M_1 \cdot G_i \cdot M_2 = \{M_1 X M_2 | X \in G_i\}$ results in a new set of size q^d whose pairwise differences are all invertible but is not multiplicatively closed in general. Although the field embedding technique is perhaps the most elegant way of building G_i, and certainly the most constructive, it may transpire that taking s from some set with less algebraic structure is advantageous in terms of the hardness of the resulting search problem. One can also construct the valid set $G_i + X$ by adding a fixed matrix X to each element of G_i, but this technique is somewhat constrained by the fact that LWE samples are additive in the secret s (e.g., one could just add $a \cdot X$ into the second coordinate of the resulting samples).

Although this restriction is not ideal, we have a remark about the implications on the security of the CLWE problem. Restricting the secret space in (R)LWE problems is not an uncommon idea: tertiary secrets, where each coordinate of $s \in \{-1, 0, 1\}$, are used in the NIST candidate LAC [24] amongst others, and security whilst restricting the secret to orders or subfields is discussed in [11], and to other K-lattices in [39]. Overall, we suspect that the decision CLWE problem is polynomial time equivalent to the search CLWE problem without restriction on s, in particular when the number of samples is small as in our applications in Sect. 5, and that the restriction is a function of our reduction technique rather than some causal property of the CLWE distribution. For the purposes of constructing a cryptosystem, we assume that this reduction implies that the decision CLWE problem is hard.

[8] Although this case exists only when each $\mathfrak{q}_i \mathcal{O}_L$ is a prime ideal in $\mathcal{O}_L$.

E. The Case where q Totally Ramifies in Relative Extension L/K

Here, we apply a decomposition in terms of Λ ideals:

$$\Lambda_q = \Lambda/q\Lambda = \Lambda/\mathcal{P}_1^{e_1}...\mathcal{P}_g^{e_g}, \tag{4}$$

where the $\mathcal{P}_i$ are maximal two-sided ideals in Λ and the e_i are some positive integers. Moreover, the following holds (see [30]):

$$\Lambda/\mathcal{P}_i \cong M_{f_i}(\mathbb{F}_{q^{e_i}}),$$

where $f_i e_i = d$. When $e_i = d$, we have $\Lambda/\mathcal{P}_i \cong M_1(\mathbb{F}_{q^d}) = \mathbb{F}_{q^d}$, a finite field.

We reduce CLWE to CLWE modulo $\mathcal{P}_i^d$ using a similar proof as above, and from there reduce to CLWE modulo $\mathcal{P}_i$. The secret then lies in some finite field, so the difference of any two elements will invert and the size of the secret space will be unrestricted. However, in order to achieve this we will have to consider the reduction for ideal lattice problems where the ideal is coprime to the ideal generated by the modulus q. This is still an infinite set of ideal lattices. Before proceeding with the reduction, we first remove the restriction on the ramification of the modulus present in the statements of the technical lemmas.

In [34], Propositions 1 and 4 state that for $\mathfrak{p}_i \subset \mathcal{O}_K$ unramified, and inert or split in $\mathcal{O}_L$, $\mathfrak{p}_i\Lambda = \bigoplus_{j=0}^{d-1} u^j \mathfrak{p}_i \mathcal{O}_L$, and the $\mathfrak{p}_i\Lambda$ are the largest two-sided ideals containing $q\Lambda$. In our case, we are dealing with $\mathfrak{p}_i$ ramified and not split in $\mathcal{O}_L$.

Let $p \in \mathbb{Z}$ be a prime such that $p\mathcal{O}_K = \mathfrak{p}_1...\mathfrak{p}_{[K:\mathbb{Q}]}$. Moreover, let $\mathfrak{p}_i\mathcal{O}_L = (\mathfrak{P}_1...\mathfrak{P}_g)^e$, where $eg = [L : K]$ and $e > 1$; importantly, this means that $f_{\mathfrak{P}_i} = [\mathcal{O}_L/\mathfrak{P}_i : \mathcal{O}_K/\mathfrak{p}_i] = 1$. Set $\mathcal{I} = \mathfrak{P}_1...\mathfrak{P}_g \oplus u\mathfrak{P}_1...\mathfrak{P}_g \oplus ... \oplus u^{d-1}\mathfrak{P}_1...\mathfrak{P}_g$ in $\Lambda = \mathcal{O}_L \oplus u\mathcal{O}_L \oplus ... \oplus u^{d-1}\mathcal{O}_L$. It can be verified that $\mathcal{I}$ is a two-sided ideal.

Background on the following definitions can be found in [43].

Definition 28. The order ideal $ord_{\mathcal{O}_K}(X)$ of a finitely generated $\mathcal{O}_K$-module X is defined as follows:

1. If $X = 0$, $ord_{\mathcal{O}_K}(X) = \mathcal{O}_K$;
2. If X is not an $\mathcal{O}_K$-torsion module, $ord_{\mathcal{O}_K}(X) = 0$;
3. If X is a nonzero $\mathcal{O}_K$-torsion module, then X has an $\mathcal{O}_K$-composition series, whose composition factors are $\{\mathcal{O}_K/\mathfrak{p}_i\}$, with $\mathfrak{p}_i$ ranging over some set of maximal ideals of $\mathcal{O}_K$. Set $ord_{\mathcal{O}_K}(X) = \prod_i \mathfrak{p}_i$, where the number of factors equals the number of composition factors of X.

Definition 29. Let M be an integral ideal of Λ. Define its *norm* by

$$N_{\mathcal{A}/K}(M) = ord_{\mathcal{O}_K}\Lambda/M$$

Lemma 23. (24.6 of [43]) *Let $\mathcal{J}$ be a prime ideal of Λ, and let $\mathcal{J} \cap \mathcal{O}_K = \mathfrak{p}$. Set $f = [\Lambda/\mathcal{J} : \mathcal{O}_K/\mathfrak{p}]$. Then $N_{\mathcal{A}/K}(\mathcal{J}) = \mathfrak{p}^f$.*

Lemma 24. (Theorem 24.13 of [43]) *For any maximal integral ideal M, $N_{rd}(M) = \mathfrak{p}$ for M lying above $\mathfrak{p}$, if $\mathcal{O}_K/\mathfrak{p}$ is a finite field.*

To prove the desired result we use a norm argument, considering the norm of $\mathcal{I}$, $N_{\mathcal{A}/K}(\mathcal{I})$, defined in Definition 29 to be $ord_{\mathcal{O}_K}(\Lambda/\mathcal{I})$. What is $ord_{\mathcal{O}_K}(\Lambda/\mathcal{I})$? Since $\Lambda/\mathcal{I}$ is non-zero, $ord_{\mathcal{O}_K}(\Lambda/\mathcal{I}) \neq \mathcal{O}_K$. Furthermore, $\Lambda/\mathcal{I}$ has $\mathcal{O}_K$-torsion: observe that $(\mathcal{O}_K \cap \mathcal{I})(x + \mathcal{I}) \subset \mathcal{I}(x + \mathcal{I}) \in \mathcal{I}$, for all $x \in \Lambda$, so $(\mathcal{O}_K \cap \mathcal{I})(\Lambda/\mathcal{I}) = 0$ and $\Lambda/\mathcal{I}$ is an $\mathcal{O}_K$-torsion module. Thus, $ord_{\mathcal{O}_K}(\Lambda/\mathcal{I}) \neq 0$. This leaves 3. Composition series can be hard to figure out explicitly, but in fact our calculation of $ord_{\mathcal{O}_K}(\Lambda/\mathcal{I})$ will reduce to figuring out $ord_{\mathcal{O}_K}(\mathcal{O}_L/\mathcal{J})$, for some ideal $\mathcal{J}$ of $\mathcal{O}_L$. This has an easy description when $\mathcal{J}$ is a product of prime ideals: $ord_{\mathcal{O}_K}(\mathcal{O}_L/\mathfrak{P}) = \mathfrak{p}^{f_{L/K}}$, where $\mathfrak{P} \cap \mathcal{O}_K = \mathfrak{p}$ and $f_{L/K} = [\mathcal{O}_L/\mathfrak{P} : \mathcal{O}_K/\mathfrak{p}]$, the inertial degree. So $ord_{\mathcal{O}_K}(\mathcal{O}_L/\mathfrak{P}) = N_{L/K}(\mathfrak{P})$ (see [43], 4.33).

Proposition 2. *Let $\mathfrak{p}\mathcal{O}_L = (\mathfrak{P}_1 ... \mathfrak{P}_g)^e$, where $eg = [L : K]$ and $e > 1$. Set $\mathcal{I} = \mathfrak{P}_1 ... \mathfrak{P}_g \oplus u\mathfrak{P}_1 ... \mathfrak{P}_g \oplus ... \oplus u^{d-1}\mathfrak{P}_1 ... \mathfrak{P}_g$. Then $\mathcal{I}$ is a maximal ideal in Λ.*

Proof. We consider two related norms, the norm from $\mathcal{A}$ to K, denoted $N_{\mathcal{A}/K}$, and the reduced norm, denoted N_{rd}. They are related as follows: $N_{\mathcal{A}/K} = N_{rd}^d$, where $[L : K] = d$. In our case the inertial degree $f_{L/K} = 1$, so $\mathcal{O}_L/\mathfrak{P}_j \cong \mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_p$, and $[\mathcal{O}_L/\mathfrak{P}_j : \mathcal{O}_K/\mathfrak{p}] = 1$. Moreover, we have

$$\begin{aligned}
\Lambda/\mathcal{I} &= (\mathcal{O}_L \oplus u\mathcal{O}_L \oplus ... \oplus u^{d-1}\mathcal{O}_L)/(\mathfrak{P}_1 ... \mathfrak{P}_g \oplus u\mathfrak{P}_1 ... \mathfrak{P}_g \oplus ... \oplus u^{d-1}\mathfrak{P}_1 ... \mathfrak{P}_g) \\
&\cong \mathcal{O}_L/\mathfrak{P}_1 ... \mathfrak{P}_g \oplus u\mathcal{O}_L/u\mathfrak{P}_1 ... \mathfrak{P}_g \oplus ... \oplus u^{d-1}\mathcal{O}_L/u^{d-1}\mathfrak{P}_1 ... \mathfrak{P}_g \\
&\cong (\mathcal{O}_L/\mathfrak{P}_1 ... \mathfrak{P}_g)^d \cong (\mathcal{O}_K/\mathfrak{p})^{g \cdot d},
\end{aligned}$$

so $f = gd$. Thus, if $\mathcal{I}$ is prime, by Lemma 23 above, $N_{\mathcal{A}/K}(\mathcal{I}) = \mathfrak{p}^{gd}$. We have:

$$\begin{aligned}
N_{\mathcal{A}/K}(\mathcal{I}) &= ord_{\mathcal{O}_K}(\Lambda/\mathcal{I}) = ord_{\mathcal{O}_K}((\mathcal{O}_L/\mathfrak{P}_1 ... \mathfrak{P}_g)^d) = ord_{\mathcal{O}_K}(\mathcal{O}_L/\mathfrak{P}_1^d ... \mathfrak{P}_g^d) \\
&= N_{L/K}(\mathfrak{P}_1^d ... \mathfrak{P}_g^d) = N_{L/K}(\mathfrak{P}_1^d) ... N_{L/K}(\mathfrak{P}_g^d) = N_{L/K}(\mathfrak{P}_1)^d ... N_{L/K}(\mathfrak{P}_g)^d \\
&= ord_{\mathcal{O}_K}(\mathcal{O}_L/\mathfrak{P}_1)^d ... ord_{\mathcal{O}_K}(\mathcal{O}_L/\mathfrak{P}_g)^d = \mathfrak{p}^d ... \mathfrak{p}^d = \mathfrak{p}^{gd},
\end{aligned}$$

as required. So $\mathcal{I}$ has the same norm as a prime ideal.

We finally show that if $\mathcal{I}$ were not a maximal two-sided ideal (so prime), then we obtain a contradiction. Suppose we have $\mathcal{I} \subsetneq \mathcal{J} \subsetneq \Lambda$, where $\mathcal{J}$ is a maximal two-sided ideal of Λ. Then $|\Lambda/\mathcal{J}| < |\Lambda/\mathcal{I}|$, and so $[\Lambda/\mathcal{J} : \mathcal{O}_K/\mathfrak{p}] < [\Lambda/\mathcal{I} : \mathcal{O}_K/\mathfrak{p}]$, or equivalently $f_{\mathcal{J}} < f_{\mathcal{I}}$ for f as defined previously. Then $N_{\mathcal{A}/K}(\mathcal{I}) = \mathfrak{p}^{f_{\mathcal{I}}} \subsetneq \mathfrak{p}^{f_{\mathcal{J}}} = N_{\mathcal{A}/K}(\mathcal{J})$; using the relation between the norms gives $N_{rd}(\mathcal{I}) \subsetneq N_{rd}(\mathcal{J})$, which are both ideals of $\mathcal{O}_K$ - but $N_{rd}(\mathcal{I})$ is maximal in $\mathcal{O}_K$, so $N_{rd}(\mathcal{J})$ cannot be a proper ideal containing it. This is a contradiction, and the result follows. □

Corollary 2. *Let $\mathfrak{p}_i \subset \mathcal{O}_K$ be a prime ideal above prime $q \in \mathbb{Z}$, such that $\mathfrak{p}_i\mathcal{O}_L = \mathfrak{P}_i^e$, for some positive integer $e \leq [L : K] = d$. Then $\mathcal{I} = \mathfrak{P}_i + u\mathfrak{P}_i + ... + u^{d-1}\mathfrak{P}_i$ is the maximal ideal of Λ lying above $\mathfrak{p}_i$.*

Proof. We have three statements to prove: that $\mathcal{I}$ is a two-sided ideal, that it is maximal, and that it lies above $\mathfrak{p}_i$. The latter statement is clear: $\mathcal{I} \cap \mathcal{O}_K = \mathfrak{P}_i \cap \mathcal{O}_K = \mathfrak{p}_i$. Moreover, maximality follows from Proposition 2.

To see that it is an ideal, first note that it is additively closed. In addition, for any element of $Gal(L/K)$, say θ, we have $\theta(\mathfrak{P}_i) = \mathfrak{P}_i$, because the automorphism permutes the primes above $\mathfrak{p}_i$, and there is only one that can be permuted. We now drop the subscript and write $\mathfrak{P}$. Let $a \in \mathcal{I}$ and $b \in \Lambda$. Then

1. $$\begin{aligned} a \cdot b &= (a_1 + ua_2 + \ldots + u^{d-1}a_{d-1}) \cdot (b_1 + ub_2 + \ldots + u^{d-1}b_{d-1}) \\ &= \sum_{j=0}^{d-1} u^j \gamma^{\alpha_{ijk}} \sum_{i+k \equiv j \bmod d}^{d-1} \theta^k(a_i) b_k \subset \sum_{j=0}^{d-1} u^j \gamma^{\alpha_{ijk}} \sum_{i+k \equiv j \bmod d}^{d-1} \theta^k(\mathfrak{P}) b_k \\ &\subset \sum_{j=0}^{d-1} u^j \gamma^{\alpha_{ijk}} \sum_{i+k \equiv j \bmod d}^{d-1} \mathfrak{P} \subset \mathfrak{P} \oplus u\mathfrak{P} \oplus \ldots \oplus u^{d-1}\mathfrak{P} = \mathcal{I}, \end{aligned}$$

2. and $$\begin{aligned} b \cdot a &= (b_1 + ub_2 + \ldots + u^{d-1}b_{d-1}) \cdot (a_1 + ua_2 + \ldots + u^{d-1}a_{d-1}) \\ &= \sum_{j=0}^{d-1} u^j \gamma^{\alpha_{ijk}} \sum_{i+k \equiv j \bmod d}^{d-1} \theta^k(b_i) a_k \subset \sum_{j=0}^{d-1} u^j \gamma^{\alpha_{ijk}} \sum_{i+k \equiv j \bmod d}^{d-1} \theta^k(b_i) \mathfrak{P} \\ &\subset \sum_{j=0}^{d-1} u^j \gamma^{\alpha_{ijk}} \sum_{i+k \equiv j \bmod d}^{d-1} \mathfrak{P} \subset \mathfrak{P} \oplus u\mathfrak{P} \oplus \ldots \oplus u^{d-1}\mathfrak{P} = \mathcal{I}, \end{aligned}$$

where $\alpha_{ijk} = \begin{cases} 1, & i + k \neq j \\ 0, & i + k = j \end{cases}$. Thus, $\mathcal{I}$ is closed by multiplication on both sides. □

We can use our result on maximal ideals to say the following:

Lemma 25. *Assume $q \in \mathbb{Z}$ is prime such that q is completely split in $\mathcal{O}_K$, $f^q_{L/\mathbb{Q}} = 1$, and $e^q_{L/K} > 1$. Let $\mathcal{I} \subset \Lambda$ be an ideal not contained in the same maximal ideal as $q\Lambda$, and let $\mathcal{J} = q \cdot \Lambda = \langle q \rangle \cdot \Lambda$, where q is a prime integer and $\langle q \rangle = \prod_{i=1}^r \mathfrak{q}_i$ is a decomposition into prime ideals in $\mathcal{O}_K$. Assume $\gamma \notin \mathfrak{q}_i$ for each i. Then, there exists an element $t \in \mathcal{I} \cap \mathcal{O}_K$ such that $t \cdot \mathcal{I}^{-1} \subset \Lambda$ is coprime to $\mathcal{J}$, and we can compute such a t efficiently given $\mathcal{I}$ and the prime factorization of $\mathcal{J}$.*

Proof. For an ideal $\mathcal{I}$ denote by $\overline{\mathcal{I}}$ its intersection with K, which is a non-trivial ideal of $\mathcal{O}_K$. As usual, we obtain $t \in \overline{\mathcal{I}}$ such that $t \cdot \overline{\mathcal{I}}^{-1}$ and $\overline{\mathcal{J}}$ are coprime as ideals of $\mathcal{O}_K$ and $t \in \overline{\mathcal{I}} \setminus \bigcup_{i=1}^r \mathfrak{q}_i \cdot \overline{\mathcal{I}}$. Assume, for a contradiction, that $t \cdot \mathcal{I}^{-1} + \mathcal{J} \neq \Lambda$, i.e., the ideals are not coprime. Then, there is some maximal ideal $\mathcal{M}$ of Λ containing $t \cdot \mathcal{I}^{-1}$ and $\mathcal{J}$. Write $\mathfrak{q}_i \mathcal{O}_L = (\mathfrak{P}_1 \ldots \mathfrak{P}_g)^e$. Since q is has inertial degree equal to 1 in $\mathcal{O}_L$ and $\gamma \notin \mathfrak{q}_i$, by the theorem in the previous section, this ideal must be one of the form $\mathfrak{P}_1 \ldots \mathfrak{P}_g \oplus u\mathfrak{P}_1 \ldots \mathfrak{P}_g \oplus u^{d-1}\mathfrak{P}_1 \ldots \mathfrak{P}_g$ since it contains $\mathcal{J}$. Then $t \cdot \mathcal{I}^{-1} \subset \mathfrak{P}_1 \ldots \mathfrak{P}_g \oplus u\mathfrak{P}_1 \ldots \mathfrak{P}_g \oplus u^{d-1}\mathfrak{P}_1 \ldots \mathfrak{P}_g$ and consequentially $t \in (\mathfrak{P}_1 \ldots \mathfrak{P}_g \oplus u\mathfrak{P}_1 \ldots \mathfrak{P}_g \oplus u^{d-1}\mathfrak{P}_1 \ldots \mathfrak{P}_g) \cdot \mathcal{I}$ because $\mathcal{I} \cdot \mathcal{I}^{-1} = \Lambda$ in a maximal order. Since t is central it follows that $t \in ((\mathfrak{P}_1 \ldots \mathfrak{P}_g \oplus u\mathfrak{P}_1 \ldots \mathfrak{P}_g \oplus u^{d-1}\mathfrak{P}_1 \ldots \mathfrak{P}_g) \cdot \mathcal{I}) \cap \mathcal{O}_K$. Thus, we have $t \in \mathfrak{q}_i$ and $t \in \overline{\mathcal{I}}$, i.e., $t \in \mathfrak{q}_i \cap \overline{\mathcal{I}}$. Since $\mathcal{I}$ is not contained in any of the maximal ideals above q, $\overline{\mathcal{I}}$ lies above an integer m where $gcd(q, m) = 1$. Bezout's theorem tells us that there exist $a, b \in \mathbb{Z}$ such that $aq + bm = 1$. Thus, $\mathfrak{q}_i$ and $\overline{\mathcal{I}}$ are coprime, and $t \in \mathfrak{q}_i \cap \overline{\mathcal{I}} = \mathfrak{q}_i \overline{\mathcal{I}}$—which is a contradiction. □

Note here we have had to impose an extra condition—that $\mathcal{I}$ does not share a maximal ideal with q. This means that the relevant intersections with $\mathcal{O}_K$ are coprime ideals, and the proof goes through. This is not a particularly strong restriction, as there are many such ideals $\mathcal{I}$.

Lemma 26. *Let Λ, γ, and q be given in Lemma 3. Let $\mathcal{I}, \mathcal{J}$ be ideals of Λ as above, with $t \in \mathcal{I} \cap \mathcal{O}_K$ chosen as above such that $t \cdot \mathcal{I}^{-1}$ and $\mathcal{J}$ are coprime as ideals, and let $\mathcal{P}$ denote an arbitrary fractional ideal of Λ. Then, the function $\chi_t : \mathcal{A} \to \mathcal{A}$ defined as $\chi_t(x) = t \cdot x$ induces a module isomorphism from $\mathcal{P}/\mathcal{J} \cdot \mathcal{P} \to \mathcal{I} \cdot \mathcal{P}/\mathcal{I} \cdot \mathcal{J} \cdot \mathcal{P}$. Furthermore, in the case $\mathcal{J} = \langle q \rangle$ for a prime integer q we can efficiently compute the inverse.*

Proof. The proof only relies on the ramification of q insofar as the above lemma does, so the proof holds under the conditions of the previous lemma. □

The above results mean that, subject to the weak condition in Lemma 25, the reduction to search CLWE in the main body of the paper holds for primes q such that q is split completely in $\mathcal{O}_K$, and has $f^q_{L/\mathbb{Q}} = 1$, using an ideal $\mathcal{I} \in \Lambda$ that doesn't share a maximal ideal with the prime q. This removes the restrictions on q, and we have traded q unramified in $\mathcal{O}_L$ with arbitrary ideal $\mathcal{I}$, for q having $f^q_{L/\mathbb{Q}} = 1$ with $\mathcal{I}$ containing any integer which is coprime to q. There has been a tradeoff between the number of valid primes and the number of valid ideals.

The following is the first step in the reduction using ramified primes.

Reducing CLWE to CLWE modulo $\mathcal{P}_i^d$ As above, we use the extended embeddings of K to $\mathcal{A}$. Since any embedding of K can be extended to an embedding of L, we use those extended embeddings to send $\mathcal{A} = (L/K, \theta, \gamma)$ to $\mathcal{A}' = (L/K, \theta, \gamma')$, where γ' is the image of γ under a chosen embedding. These maps preserve the decomposition of $\Lambda_q^\vee$ by sending $\mathcal{P}_i$ to some $\mathcal{P}_j$—we below show that these embeddings permute the primes $\mathcal{P}_i$ modulo $q\Lambda$. We will abuse notation and denote the action of α on the cosets $\Lambda/q\Lambda$ also by α.

Lemma 27. *Let α be an isomorphism from $\mathcal{A} \to \mathcal{A}'$ as above. Fix a prime $q \in \mathbb{Z}$ such that $q\mathcal{O}_K = \mathfrak{p}_1 ... \mathfrak{p}_g$. Let $\mathcal{P}_i$ be a prime ideal of Λ lying above the prime ideal $\mathfrak{p}_i \subset \mathcal{O}_K$. Then, considering α as acting on the cosets of $\Lambda/q\Lambda$, $\alpha(\mathcal{P}_i + q\Lambda) = \mathcal{P}_j + q\Lambda$, for some $i \neq j$.*

Proof. First observe that α permutes the primes of $\mathcal{O}_K$, since it was induced by an element of $Gal(K/\mathbb{Q})$. Thus, $\alpha(\mathfrak{p}_i) = \mathfrak{p}_j$, and so $\alpha(\mathfrak{p}_i + q\Lambda) = \mathfrak{p}_j + q\Lambda$, where we have used that α fixes Λ_q. Moreover, since $\mathfrak{p}_i \subset \mathcal{P}_i$, we have $\mathfrak{p}_j + q\Lambda = \alpha(\mathfrak{p}_i + q\Lambda) \subset \alpha(\mathcal{P}_i + q\Lambda) = \alpha(\mathcal{P}_i) + q\Lambda$. Since α fixes $\Lambda/q\Lambda$, we in fact have that $\alpha(\mathcal{P}_i) + q\Lambda \subset \Lambda/q\Lambda$. Note that $\alpha(\mathcal{P}_i)$ is a prime (and hence maximal) $\alpha(\Lambda)$ ideal. Thus, $\alpha(\mathcal{P}_i) + \alpha(q\Lambda) = \alpha(\mathcal{P}_i) + q\Lambda$ is a prime ideal of $\alpha(\Lambda)/\alpha(q\Lambda) = \alpha(\Lambda/q\Lambda) = \Lambda/q\Lambda$. So we find that $\alpha(\mathcal{P}_i + q\Lambda)$ corresponds to a maximal ideal of $\Lambda/q\Lambda$ lying above $\mathfrak{p}_j$; thus, $\alpha(\mathcal{P}_i + q\Lambda) = \mathcal{P}_j + q\Lambda$. □

Lemma 28. (Reduction from CLWE to $\mathcal{P}_i^d$-CLWE) *There is a deterministic polynomial time reduction from CLWE$_{q,\Sigma_\alpha}$ to $\mathcal{P}_i^d$-CLWE$_{q,\Sigma_\alpha}$.*

Proof. Let $\mathcal{O}_i$ denote an oracle for the $\mathcal{P}_i^d$-CLWE$_{q,\Sigma}$ problem. Equation (4) defines an isomorphism, so we can use the oracle $\mathcal{O}_i$ to solve the $\mathcal{P}_j^d$-CLWE$_{q,\Sigma}$ problem for each j. Let $\alpha_{j/i}$ be an extension of the automorphism of K that maps $\mathfrak{q}_j$ to $\mathfrak{q}_i$.

Given sample $(a, b) \leftarrow \Pi_{q,s,\Sigma_\alpha}$, construct a sample of the form $\big(\alpha_{j/i}(a), \alpha_{j/i}(b)\big)$. Since Λ_q and $\Lambda_q^\vee$ are fixed by each $\alpha_{j/i}$, the sample is a valid CLWE sample in $\mathcal{A}' = (L/K, \theta, \alpha_{j/i}(\gamma))$. Feeding this sample into $\mathcal{O}_i$ outputs a value $t_j \bmod \mathcal{P}_i^d$.

We show that $\alpha_{j/i}^{-1}\left(t_j\right) = s \bmod \mathcal{P}_j^d$. Since $\alpha_{j/i}$ is an automorphism, each (a, b) is mapped to CLWE sample $\big(\alpha_{j/i}(a), \alpha_{j/i}(a \cdot s/q + e) \bmod \Lambda^\vee\big)$ in the algebra $\mathcal{A}'$, and we can write $\alpha_{j/i}(a) \cdot \alpha_{j/i}(s)/q + \alpha_{j/i}(e) \bmod \Lambda^\vee$. As stated above, our automorphisms fix our family of error distributions, and map the uniform distribution to the uniform distribution, so this is a valid CLWE instance with secret $\alpha_{j/i}(s) \in \alpha_{j/i}(\Lambda_q^\vee) = \Lambda_q^\vee$ and error distribution $\Sigma' \in \Sigma_\alpha$. So $\mathcal{O}_i$ outputs $t = \alpha_{j/i}(s) \bmod \mathcal{P}_i^d$, which yields $\alpha_{j/i}^{-1}(t) = s \bmod \mathcal{P}_j^d$, since the embeddings permute the $\mathcal{P}_i$, and thus the $\mathcal{P}_i^d$, as required. □

CLWE Modulo $\mathcal{P}_i$ We now show that it suffices to solve the problem modulo $\mathcal{P}_i$, rather than modulo $\mathcal{P}_i^d$. Since s is not zero in $\Lambda/q\Lambda$, s is not in $\mathcal{P}_1^d \ldots \mathcal{P}_g^d$, so there exists a $k : s \notin \mathcal{P}_k^d$. We will first show that the corresponding problem for RLWE can be solved; we will then show that the problem for CLWE can be solved using the method for RLWE. First, we need some lemmas and definitions.

RLWE Let $R = \mathbb{Z}[x]/\Phi_n(x)$, where $\Phi_n(x)$ is the nth cyclotomic polynomial. Then R is the ring of integers of the nth cyclotomic field, denoted K. Let $R_p = R/pR$, and $R^\vee = \{x \in K : Tr(xR) \subset \mathbb{Z}\}$ be the dual lattice. An RLWE sample has the form $(a, b) = (a, (a \cdot s)/p + e \bmod R^\vee) \in R_p \times \mathbb{T}$, where $a \leftarrow R_p$ uniformly at random, $s \leftarrow R_p^\vee$, and e sampled according to some error distribution; finally, $\mathbb{T}$ is the unit torus. Let $p\mathcal{O}_K = \prod_{i=0}^r \mathfrak{p}_i^e$, for $e > 1$. Let $\mathfrak{p}_{i,j}$-RLWE be the problem of finding $s \bmod \mathfrak{p}_i^j$, given RLWE sample (a, b). We show that we can solve this problem, given access to a $\mathfrak{p}_{i,1}$-RLWE oracle. Note that knowing $s \bmod \mathfrak{p}_i^e$ is sufficient to find s, by using automorphisms and the CRT.

Lemma 29. *Given RLWE sample (a, b) and an oracle for $\mathfrak{p}_{i,1}$-RLWE oracle, we can solve $\mathfrak{p}_{i,e}$-RLWE.*

Proof. Let (a, b) be an RLWE sample, and submit (a, b) to the oracle to obtain an element x such that $x \equiv s \bmod \mathfrak{p}_i$. Then $x - s \in \mathfrak{p}_i$. We can write $x - s = \alpha \cdot p + \beta \cdot f_i(\zeta_n)$, where $\alpha, \beta \in \mathcal{O}_K$, since $\mathcal{O}_K = \mathbb{Z}[\zeta_n]$ for cyclotomic fields, and $\mathfrak{p} = (p, f_i(\zeta_n))$, where $\Phi_n(x) = \prod_{i=0} f_j(x) \bmod p$, and the f_j are irreducible modulo p. So $x - s - \beta f_i(\zeta_n) = \alpha \cdot p \subset p\mathcal{O}_K$, and $s \equiv x - \beta f_i(\zeta_n) \bmod p\mathcal{O}_K$. We proceed to construct an element congruent to s modulo $p^e\mathcal{O}_K$, since $p^e\mathcal{O}_K \subset \mathfrak{p}_i^e\mathcal{O}_K$.

Replace s by $(x - s - \beta f_i(\zeta_n))/p$. Then $(a, b') = (a, (a \cdot (\frac{1}{p}(x - s - \beta f_i(\zeta_n)))/p + e' \bmod R^\vee)$ is a valid RLWE sample. Submit it to the oracle to obtain y such that

$y \equiv \frac{x-s-\beta f_i(\zeta_n)}{p} \bmod \mathfrak{p}_i$. As before, write this in terms of the generators of $\mathfrak{p}$, and subtract the $f_i(\zeta_n)$ term to obtain an element in $p\mathcal{O}_K$, resulting in $y - d \cdot f_i(\zeta_n) - \frac{x-s-\beta f_i(\zeta_n)}{p} = p \cdot e$ for some d and $e \in \mathcal{O}_k$. Replace $\frac{x-s-\beta f_i(\zeta_n)}{p}$ by $\frac{y-d\cdot f_i(\zeta_n)}{p} - \frac{x-s-\beta f_i(\zeta_n)}{p^2}$. Continue in this manner until we have $\frac{v-w\cdot f_i(\zeta_n)}{p} - \ldots - \frac{y-d\cdot f_i(\zeta_n)}{p^{e-1}} - \frac{x-s-\beta f_i(\zeta_n)}{p^e} = z \in \mathcal{O}_K$. Rearrange for $s = p^e z - p^{e-1}(v - w \cdot f_i(\zeta_n)) + \ldots + p(y - d \cdot f_i(\zeta_n)) + x - \beta f_i(\zeta_n)$. Clearly $s \equiv p^e z - p^{e-1}(v - w \cdot f_i(\zeta_n)) + \ldots + p(y - d \cdot f_i(\zeta_n)) + x - \beta f_i(\zeta_n) \equiv x \bmod \mathfrak{p}_i$. Moreover by construction we have found an element in the same coset modulo p^e as s, namely $p^{e-1}(v - w \cdot f_i(\zeta_n)) - \ldots - p(y - d \cdot f_i(\zeta_n)) - x + \beta f_i(\zeta_n)$. Reducing modulo $\mathfrak{p}_i^e$, we obtain an element of $\mathcal{O}_K$ congruent to s, which is a solution to $\mathfrak{p}_{i,e}$-RLWE. □

Solving $\mathcal{P}_{i,d}$-CLWE

Lemma 30. *Let $q \in \mathbb{Z}$ be prime such $q\mathcal{O}_K = \prod \mathfrak{p}_i$, $\mathfrak{p}_i\mathcal{O}_L = \mathfrak{P}_i^d$ and $q\Lambda = \mathcal{P}_1^d \ldots \mathcal{P}_g^d$. Given CLWE sample (a, b) and an oracle for the CLWE* mod $\mathcal{P}_i$ *problem, we can solve the CLWE* mod$\mathcal{P}_i^d$*-problem.*

Proof. Submit (a, b) to the oracle for $x \in \Lambda_q^\vee$: $x \equiv s \bmod \mathcal{P}_i$. By Proposition 2, $\mathcal{I} = \mathfrak{P}_i + u\mathfrak{P}_i + \ldots + u^{d-1}\mathfrak{P}_i$ is the maximal ideal of Λ lying above $\mathfrak{p}_i$. So we can take $\mathcal{P}_i = \mathfrak{P}_i + u\mathfrak{P}_i + \ldots + u^{d-1}\mathfrak{P}_i$. Then $x - s \in \mathcal{P}_i$, and hence $x_i - s_i \in \mathfrak{P}_i$ for each i, where x_i and s_i are the ith coefficient of x and s respectively.
The prime ideals of the ring of integers of an algebraic number field lying above the prime q have the form $\mathfrak{P}_i = (q, f_i(\alpha))$, for some polynomial f_i and $\alpha \in \mathcal{O}_L$. Thus, proceeding as in the RLWE case, we can express $x_i - s_i$ in terms of q and $f_i(\alpha)$, subtract the $f_i(\alpha)$ term, and have an element divisible by q. We replace the s_i with the resulting element, $\frac{x_i - s_i - b_i \cdot f_j(\alpha)}{q}$, for each i, to obtain a new valid CLWE sample with new secret x', and then query the oracle for a value congruent to x' modulo $\mathcal{P}_i$. We can iterate the procedure as before, until we have an element y_i such that $y_i \equiv s_i \bmod \mathfrak{P}_i^d$.

We can then obtain an element y such that $y_i - s_i$ is divisible by q^d for each i, and hence $y - s$ is divisible by q^d, so $y - s \in \mathcal{P}_i^d$. □

This lemma means that if we can solve search CLWE modulo $\mathcal{P}_i$, we can construct a solution to search CLWE modulo $\mathcal{P}_i^d$; we can then use the argument of the preceding section (using the embeddings and the CRT) to find the secret s and solve CLWE.
In the following section, in a series of steps mirroring the standard methods, adapted largely from [26], we establish the hardness of the decision problem.

Hybrid CLWE and Search to Decision

Definition 30. For $s \in \Lambda_q^\vee$, distribution Σ over $\oplus_j u^j L_\mathbb{R}$, and $i \in [n]$, define a sample from distribution $\Pi_{q,s,\Sigma}^i$ over $\Lambda_q \times \left(\oplus_{j=0}^{d-1} u^j L_\mathbb{R}\right)/\Lambda^\vee$ by taking $(a, b) \leftarrow \Pi_{q,s,\Sigma}$ and $h \in \Lambda_q^\vee$ which is uniformly random and independent mod $\mathcal{P}_j$, for $j \leq i$ and 0 mod $\mathcal{P}_j$, for $j > i$, and outputting $(a, b + h/q)$. If $i = 0$, define $\Pi_{q,s,\Sigma}^0 = \Pi_{q,s,\Sigma}$. Then for

$i \in [n]$ and a family of distributions Σ_α, the WD-CLWE$^i_{q,\Sigma_\alpha}$ problem is to find j given access to $\Pi^j_{q,s,\Sigma}$ for $j \in \{i-1, i\}$ and CLWE secret and error distribution s, Σ.

Lemma 31. *For any $i \in [n]$ there is a probabilistic polynomial-time reduction from $\mathcal{P}_i$-CLWE$_{q,s,\Sigma_\alpha,G}$ to WD-CLWE$^i_{q,s,\Sigma_\alpha}$.*

Proof. We proceed as usual. There are $|\Lambda/\mathcal{P}_i|$ possible values of s mod $\mathcal{P}_i$, which is bounded above by $|\Lambda/\mathcal{P}_i| = q^d$, so we may efficiently enumerate over the possible values. We want a transform which takes $g \in \Lambda/\mathcal{P}_i$ and maps $\Pi_{q,s,\Sigma}$ to $\Pi^{i-1}_{q,s,\Sigma}$ if $g = s$ mod $\mathcal{P}_i$ or to $\Pi^i_{q,s,\Sigma}$ otherwise. Take CLWE sample $(a, b) \leftarrow \Pi_{q,s,\Sigma}$, and output

$$\left(a', b'\right) = (a + v, b + (h + vg)/q) \in \Lambda_q \times \left(\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}\right) / \Lambda^\vee,$$

with $v \in \Lambda_q$ uniformly random mod $\mathcal{P}_i$ and 0 mod $\mathcal{P}_j$ for $j \neq i$, and $h \in \Lambda_q^\vee$ uniformly random and independent mod $\mathcal{P}_j$ for $j < i$ and 0 on the other $\mathcal{P}_j$. Then a' is uniformly distributed on Λ_q, so it remains to prove b' is distributed correctly. Fix a', then

$$\begin{aligned} b' &= b + (h + vg)/q \\ &= (as + h + vg)/q + e \\ &= \left(a's + h + v(g - s)\right)/q + e \end{aligned}$$

where e is drawn from Σ. If $g = s$ mod $\mathcal{P}_i$, then $v(g-s) = 0$ mod $\mathcal{P}_i$ so the distribution of $\left(a', b'\right)$ is $\Pi^{i-1}_{q,s,\Sigma}$. Otherwise, $v(g - s)$ is uniformly random mod $\mathcal{P}_i$ (since $\Lambda/\mathcal{P}_i$ is a field) and 0 modulo the other $\mathcal{P}_j$. Setting $h' = h + v(g - s)$, one can see that the distribution of $\left(a', b'\right)$ is $\Pi^i_{q,s,\Sigma}$, as required. □

Worst-Case to Average-Case Decision Reduction

This stage of the reduction holds identically to that of the main body of the paper, replacing $\mathcal{R}_i$ with $\mathcal{P}_i$.

F. Estimating the Multiplication Complexity

The overall flow to compute the multiplication is depicted in Fig. 5, which is explained in detail in the sequel.

F.1. *Algorithm for Multiplication in Cyclic Algebras*

We recall some details necessary to understand our multiplication algorithm. Recall that in the explicit constructions of Theorem 2 the base field K is cyclotomic and q is a prime integer chosen so that $\langle q \rangle$ splits completely in $\mathcal{O}_K$ as $\langle q \rangle = \mathfrak{q}_1 \ldots \mathfrak{q}_n$, where n is the dimension of K as an extension of $\mathbb{Q}$. Furthermore, the degree of L over K is a typically

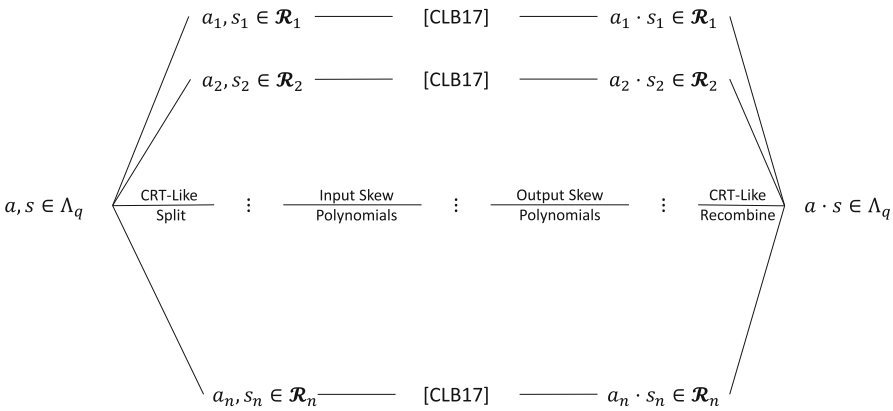

Fig. 5. Depiction of the multiplication algorithm for cyclic algebras. [CLB17] is referred to as [15].

small d. Then, following the CRT-like decomposition of Lemma 12 we write

$$\Lambda_q \cong \mathcal{R}_1 \times \cdots \times \mathcal{R}_n$$

for $\mathcal{R}_i = \bigoplus_{j=0}^{d-1} u^j \mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$. We will show that each $\mathcal{R}_i$ is a skew polynomial ring over $\mathbb{Z}_q$, and in particular a skew polynomial ring for which we can apply the algorithms of [15] to compute multiplication independently in each $\mathcal{R}_i$ in $\tilde{O}(d^\omega)$ operations in $\mathbb{Z}_q$, which output elements whose u coordinates are in the form $\sum_i \ell_i k_i$ for $k_i \in \mathcal{O}_{Kq}$ and $\{\ell_i\}$ some arbitrary normal basis for $\mathcal{O}_{Lq}$ over $\mathcal{O}_{Kq}$. We remark that the representation as a skew polynomial ring need not contradict the fact that we viewed the rings $\mathcal{R}_i$ as matrix rings in Sect. 4.2, since computing matrix multiplication can be reduced to the problem of computing multiplication of skew polynomials (see [15]). Since $\omega \leq 2.373$, this leads to a complexity of approximately $\tilde{O}(Nd^{0.373})$ and it is possible to compute the multiplication in each $\mathcal{R}_i$ in parallel. However, we must also compute the complexity of the splitting isomorphism.

F.2. *The Rings $\mathcal{R}_i$*

In order to apply the algorithm of [15], we must confirm that each $\mathcal{R}_i$ satisfies the following conditions:

- $\mathcal{R}_i$ is the quotient of a skew polynomial ring with center $\mathcal{O}_K/\mathfrak{q}_i$ by a polynomial in the form $X^d - \gamma$.
- γ is a norm from $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ into $\mathcal{O}_K/\mathfrak{q}_i$.[9]
- $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ is a field extension of $\mathcal{O}_K/\mathfrak{q}_i$ or an étale-$\mathcal{O}_K/\mathfrak{q}_i$ algebra.

The first of the conditions follows immediately from the definitions of a skew polynomial ring and a cyclic algebra. The veracity of the latter conditions will depend on how the prime ideal $\mathfrak{q}_i$ of $\mathcal{O}_K$ splits in $\mathcal{O}_L$ as $\mathfrak{q}_i\mathcal{O}_L$. Since $\mathfrak{q}_i$ is prime in K and L/K is Galois,

[9] Due to the modulo reduction this does not contradict the assumption that γ is not a global norm.

we know

$$\mathfrak{q}_i \mathcal{O}_L = \prod_{j=1}^{g} (\mathfrak{q}_{i,j})^e$$

for some prime ideals $\mathfrak{q}_{i,j}$ in $\mathcal{O}_L$ and integers e, g satisfying $efg = [L : K] = d$, where f denotes the inertial degree. Assuming that L is constructed as a subfield of a cyclotomic field as in [21], it is a Galois number field and it follows that each $\mathfrak{q}_i$ splits with the same e, f, and g. Furthermore, since they are coprime as ideals of $\mathcal{O}_K$, their factorizations' in L are disjoint. Thus, we are left to consider three cases.

We first consider the case where each $\mathfrak{q}_i \mathcal{O}_L$ remains prime in $\mathcal{O}_L$. It follows that $\mathcal{O}_L/\mathfrak{q}_i \mathcal{O}_L$ is a finite field, and computing the norm of $\mathfrak{q}_i \mathcal{O}_L$ indicates $\mathcal{O}_L/\mathfrak{q}_i \mathcal{O}_L \cong \mathbb{F}_{q^d}$. In this case it is easy to see that $\mathcal{O}_L/\mathfrak{q}_i \mathcal{O}_L$ is a finite field extension of $\mathcal{O}_K/\mathfrak{q}_i \cong \mathbb{F}_q$ and consequentially, because the norm map is surjective over finite field extensions, that γ is a norm. Here it is clear that the algorithms of [15] can be applied.

The second case we consider is $g = d$, $e = f = 1$. Now each $\mathfrak{q}_i \mathcal{O}_L$ splits completely in $\mathcal{O}_L$ into a product of prime ideals $\mathfrak{q}_{i,1} \ldots \mathfrak{q}_{i,d}$. By the CRT we have

$$\mathcal{O}_L/\mathfrak{q}_i \mathcal{O}_L \cong \bigotimes_{j=1}^{d} \mathcal{O}_L/\mathfrak{q}_{i,j}$$

where each $\mathcal{O}_L/\mathfrak{q}_{i,j} \cong \mathbb{F}_q$, and it follows that $\mathcal{O}_L/\mathfrak{q}_i \mathcal{O}_L$ is an étale-$\mathcal{O}_K/\mathfrak{q}_i$ algebra. We are left to show that γ is a norm, which we show via the stronger condition that the norm map in this extension is surjective. By the CRT, $\mathcal{O}_L/\mathfrak{q}_i \mathcal{O}_L$ is isomorphic to a direct product of d copies of $\mathbb{F}_q$. Since the embeddings of L cyclically permute the ideal factors of $\mathfrak{q}_i$ it follows that the relative norm of an element $(x_1, \ldots, x_d) \in \bigotimes_{j=0}^{d} \mathcal{O}_L/\mathfrak{q}_{i,j}$ is precisely $\prod_{k=1}^{d} x_k \mod q$. It is easy to see that this norm is surjective (because any $x \in \mathbb{F}_q$ is the norm of, e.g., $(1, 1, \ldots, x)$) and now once again we can apply the multiplication algorithms of [15].

Intermediate cases, where $\mathfrak{q}_i$ splits into a product of prime ideals with the same norm such that $e = 1$, $fg = d$, can be handled using a straightforward combination of these two methods.

The final case to consider is the ramified case, when $e \neq 1$. Now the factorization of $\mathfrak{q}_i \mathcal{O}_L$ contains some power $\mathfrak{p}_i^{e_i}$ of a prime $\mathcal{O}_L$ ideal $\mathfrak{p}_i$. In this case, we are not able to verify that the necessary conditions for the algorithms of [15] hold. However, we observe that the ideal $\langle q \rangle$ ramifies in $\mathcal{O}_L$ if and only if q divides the discriminant of $\mathcal{O}_L$. Since only a finite number of primes divide this discriminant, we restrict ourselves to considering the cases where q does not ramify. We emphasize that in the main cases of interest, where K is the mth cyclotomic field with m having small divisors and $[L : K]$ is small, it is particularly unlikely that the large modulus q typical in cryptography divides the discriminant of L. Indeed, when we pick L as a subfield of $K(\zeta_{q'})$ for some large prime integer q' using the techniques of [21] as in Theorem 2, it is easy to quantify which primes potentially ramify for a fixed choice of fields: either q' or the primes smaller than

or equal to the divisors of m. As an easy example, the modulus $q = 12289$ does not ramify in the example algebras given in the Sect. 3.4 achieving dimension 1024.

F.3. *Complexity of the CRT Style Isomorphism*

We have shown that we may apply the algorithms of [15] to compute the multiplication operation in each $\mathcal{R}_i$ in complexity $\tilde{O}(d^{\omega})$. We are left to consider the complexity of the isomorphism defined by Lemma 12 generating the rings $\mathcal{R}_i$. Essentially, this operation is a coordinatewise split of the u coordinates of $\Lambda_q = \bigoplus_{j=0}^{d-1} u^j \mathcal{O}_L$, where each entry is split into its mod $\mathfrak{q}_i \mathcal{O}_L$ parts. That is, the isomorphism maps

$$\sum_{j=0}^{d-1} u^j x_j \to \bigotimes_{i=1}^{n} \sum_{j=0}^{d-1} u^j (x_j \mod \mathfrak{q}_i \mathcal{O}_L).$$

Splitting one element $x_i \in \mathcal{O}_K$ can be done in time $O(n \log n)$ using the CRT algorithm of [28] when K is a cyclotomic field of dimension n. However, L is a not a cyclotomic field, but instead a small degree d cyclic extension of a cyclotomic. Furthermore, we are trying to split the elements of L modulo ideals of K extended to those of L. We do not know of an existing general, efficient way of doing this. The naive estimate for an optimal method would take time $O(nd \log nd)$, where nd is the dimension of L, but we suspect something this efficient is impossible. We have to perform d such splits, which would result in a total complexity of $O(N \log N/d)$. Note that this compares relatively closely with the $\tilde{O}(Nd^{0.3})$ claimed for the multiplication step, and since these steps are sequential rather than parallel which of them dominates the asymptotic complexity would depend on the exact relationship between n and d, but the result is an operational complexity essentially equivalent to that of the ring variant.

Of course, the discussion of the previous paragraph relies on our implausibly low estimate of $O(nd \log nd)$ complexity of the CRT split and so we do not claim such efficiency. Instead, we present techniques in the proceeding sections to work around the problem of splitting the L part modulo the K ideals in the factorization of q. Our methods are particularly efficient in the case where q splits completely in L, but can be generalized to arbitrary splitting at only a small cost.

F.4. *Fast Cryptography when q Splits Completely in L*

We consider an explicit method for implementing fast cryptography in the special case where the ideal $\langle q \rangle$ splits completely in $\mathcal{O}_L$. By construction, $\langle q \rangle = \prod_i \mathfrak{q}_i$ in $\mathcal{O}_K$, so in this case we split $\langle q \rangle = \prod_{i,j} \mathfrak{q}_{i,j}$ in $\mathcal{O}_L$, where the prime $\mathcal{O}_K$-ideals have prime decomposition in $\mathcal{O}_L$ denoted $\mathfrak{q}_i \mathcal{O}_L = \prod_{j=1}^{d} \mathfrak{q}_{i,j}$.

We recall some facts about the extension $\mathcal{O}_{Lq}$ of $\mathcal{O}_{Kq}$. It is clear that the extension is cyclic of degree d, with Galois group generated by θ. By the CRT,

$$\mathcal{O}_{Kq} \cong \prod_i \mathcal{O}_K/\mathfrak{q}_i \cong \mathbb{F}_q{}^n$$

$$\mathcal{O}_{Lq} \cong \prod_{i,j} \mathcal{O}_L/\mathfrak{q}_{i,j} \cong {\mathbb{F}_q}^{nd}$$

where operations on the finite field products are applied coordinatewise. We represent the CRT decomposition of $\mathcal{O}_{Lq}$ as $({\mathbb{F}_q}^d)^n$, where each copy of ${\mathbb{F}_q}^d$ corresponds to the extension $\prod_j \mathcal{O}_L/\mathfrak{q}_{i,j}$ of $\mathcal{O}_K/\mathfrak{q}_i$. In the finite field representation of $\prod_j \mathcal{O}_L/\mathfrak{q}_{i,j}$, the elements of $\mathcal{O}_K/\mathfrak{q}_i$ embed as elements of ${\mathbb{F}_q}^d$ with the same entry in each coordinate, e.g., $(x, x, \ldots, x)$, corresponding to scalars over $(\mathbb{F}_q)^d$, which can be seen from the following argument: for $k \in \mathcal{O}_K$, $k = x \mod \mathfrak{q}_i$ implies $k - x \in \mathfrak{q}_i$. Then it follows that $k - x \in \mathfrak{q}_{i,j}$ and thus $k = x \mod \mathfrak{q}_{i,j}$ for each j. Furthermore there is a simple, explicit, description of the action of θ in this representation: since θ cyclically shifts the ideals in the factorization of $\mathfrak{q}_i$, one can order each copy of ${\mathbb{F}_q}^d$ so that the action of θ on $({\mathbb{F}_q}^d)^n$ is a cyclical shift of the coordinates of each of the n copies of ${\mathbb{F}_q}^d$ concurrently. We exhibit this with a trivial example: set $d = 3$, $n = 2$. Then the action of θ on $({\mathbb{F}_q}^3)^2$ is

$$\theta(a_1, a_2, a_3, b_1, b_2, b_3) = (a_3, a_1, a_2, b_3, b_1, b_2).$$

A valid $\mathcal{O}_K/\mathfrak{q}_i$ basis for $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ of size d is $\mathbf{e}_1, \ldots, \mathbf{e}_d$, where $\mathbf{e}_i = (0, \ldots, 1, \ldots 0)$ denotes the ith element of the standard basis of dimension d. Furthermore, this basis is orthonormal in the sense that $\mathbf{e}_i \cdot \mathbf{e}_j = \mathbf{e}_i$ for $i = j$ and 0 otherwise and cyclic[10] in the sense that $\theta(\mathbf{e}_i) = \mathbf{e}_{i+1}$ (e.g., normal), since the Galois group $\langle\theta\rangle$ of L over K permutes the factors $\mathfrak{q}_{i,j}$ of $\mathfrak{q}_i\mathcal{O}_L$ for each i. Because the CRT splits $\mathcal{O}_{Lq}$ into a direct product within which operations are computed coordinatewise, we can extend this to a basis of $\mathcal{O}_{Lq}$ over $\mathcal{O}_{Kq}$ in the finite field representation by concatenating n copies of this basis together, denoting by $\mathbf{e}_i^n$ the vector of dimension nd $(\mathbf{e}_i, \mathbf{e}_i, \ldots, \mathbf{e}_i)$. This basis is still cyclic, with θ operating independently on each of the n copies of ${\mathbb{F}_q}^d$ and hence the n copies of $\mathbf{e}_i$. Concatenating the bases in this way also preserves the orthonormal property.

Denote the above basis by $\ell_1, \ldots, \ell_d$. Recall that the CRT-like decomposition Lemma 12 splits each u coordinate, an element of $\mathcal{O}_{Lq}$, into its mod $\mathfrak{q}_i\mathcal{O}_L$ parts. However, we already know the mod $\mathfrak{q}_i\mathcal{O}_L$ parts of each ℓ_j by construction. So, if we store elements of $\mathcal{O}_{Lq}$ as $\ell = \sum_{j=1}^d \ell_j k_j$ for $k_j \in \mathcal{O}_{Kq}$ we can split ℓ into its $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ components in time $O(d \cdot n \log n)$ as long as the k_j elements are stored in the polynomial representation of $\mathcal{O}_{Kq}$. Consequentially, we can perform the CRT style decomposition of an element in Λ_q whose u coordinates are stored in this manner in time $O(d^2 \cdot n \log n) = O(N \log(N/d^2))$.

Now we see a way to achieve fast multiplication in Λ_q. We are required to perform the CRT in each of the d u coordinates, after which we can plug the rings $\mathcal{R}_i$ into the fast multiplication algorithm of [15]. Since the CRT is an isomorphism and we know the image of ℓ_i under the CRT, this reduces to d copies of the CRT in $\mathcal{O}_K$, each with complexity $O(dn \log n)$, and therefore a total multiplication complexity of $O(N \log(N/d^2)) + \tilde{O}(Nd^{\omega-2})$. However, this algorithm comes with complications

[10] As long as we choose the ordering in the right way.

associated with the chosen representation of elements of $\mathcal{O}_{Lq}$, which we handle in the next section.

F.4.1. *Handling Elements in the Representation*

To use the above multiplication algorithms in the scheme of Sect. 5.2, we need to be able to store the elements compactly and sample the elements efficiently. Storing elements in this form turns out to be straightforward: each $\mathcal{O}_{Lq}$ element requires storing d elements of $\mathcal{O}_{Kq}$. An element of Λ_q is d elements of $\mathcal{O}_{Lq}$, so in total we store d^2 elements of $\mathcal{O}_{Kq}$, corresponding to one element of dimension $N = nd^2$, which is equivalent to storing d elements of dimension nd.

We now discuss how to efficiently sample elements of Λ_q according to an appropriate error distribution. Recall from the security reduction of Sect. 3 that the error distributions we recommend in practice are spherical or elliptical Gaussians in the coordinates of the embedding $\sigma_{\mathcal{A}}$. We sample using the following result.

Theorem 11. *Let L/K be a tower of number fields with $[K : \mathbb{Q}] = n$ and $[L : K] = d$ where K is a prime-power cyclotomic field. Let $q \geq 2$ be a prime modulus which splits completely in $\mathcal{O}_L$ and let $\ell_1, \ldots, \ell_d$ be the cyclic basis of $\mathcal{O}_{Lq}$ over $\mathcal{O}_{Kq}$ satisfying $\ell_i \cdot \ell_j = \ell_i$ if $i = j$ and 0 otherwise. Then, the distribution on $\mathcal{O}_{Lq}$ obtained by sampling $k_1, \ldots, k_d$ independently from a discrete Gaussian over $\mathcal{O}_{Kq}$ in the polynomial representation and outputting $\ell = \sum_i \ell_i k_i$ is a discrete Gaussian over $\mathcal{O}_{Lq}$ in the ℓ_2 norm over $L_{\mathbb{R}}$.*

Proof. Recall that in the case where K is a prime power cyclotomic the power basis is a rotation and a scaling of the canonical basis (see, e.g., [19]), so a discrete Gaussian in the polynomial representation corresponds to a discrete Gaussian in the canonical basis as well. Order the canonical embedding of $\mathcal{O}_L$ such that elements of $\mathcal{O}_K$ embed as vectors of n blocks of length d that are the same in each block, e.g.,

$$k_1 = (k_{1,1}, k_{1,1} \ldots, k_{1,1}, k_{1,2}, \ldots, k_{1,n}),$$

where each entry $k_{i,j}$ of k_i appears d times. Since the ℓ_i form a cyclic basis, in each d-block the entries of ℓ_{i+1} are just a cyclic shift of those of ℓ_i. [11] For a fixed choice of basis the distribution in each d-block of ℓ is independent, because the $k_{i,j}$ are sampled independently from a spherical Gaussian. So we can consider one d block of ℓ at a time, and write the d-block of ℓ_1 as $a_1, \ldots, a_d$. Since multiplication in the canonical embedding is coordinatewise and the ℓ_i form a cyclic basis, the first block of ℓ can be written as

$$\begin{pmatrix} a_1 & a_2 & \ldots & a_d \\ a_d & a_1 & \ldots & a_{d-1} \\ \vdots & \vdots & \ddots & \vdots \\ a_2 & a_3 & \ldots & a_1 \end{pmatrix} \cdot \begin{pmatrix} k_{1,1} \\ k_{2,1} \\ \vdots \\ k_{d,1} \end{pmatrix}.$$

[11] Again assuming a sensible ordering.

Call the left matrix $\mathbf{A}$ and the right vector $\mathbf{k}$. $\mathbf{k}$ is a Gaussian of parameter r, so $\mathbf{Ak}$ has a Gaussian distribution with covariance matrix $r \cdot \mathbf{AA}^\dagger$ by, e.g., [25, Lemma 2.5], and if this is diagonal and constant on the lead diagonal then we are done. Due to the structure of the canonical embedding and how we picked our basis in the $\mathcal{O}_L/\langle q \rangle$ representation, we have that $a_i = \theta^i(a_1)$, and that for $i \neq j$ $\theta^i(a_1) \cdot \theta^j(a_1) = 0 \mod q$. It follows that the off-diagonal entries of $\mathbf{AA}^\dagger$ are 0 (since product being 0 is preserved under representations) and the diagonal entries are $\sum_{i=1}^{d} |a_i|^2$, where $|\cdot|$ denotes the absolute value. Hence, the first d-block of ℓ is a spherical Gaussian distribution, and since this analysis holds for any block it follows that each block of ℓ is a spherical Gaussian. One also needs to show that the Gaussian distribution has the same variance in each block, but this follows from the fact that the K-embeddings permute the mod $\mathfrak{q}_i$ values and fix the ℓ_2 norm of $K_\mathbb{R}$. Explicitly, by construction each K embedding modulo $\langle q \rangle$ can be extended 'identically' onto $\mathcal{O}_L \mod \langle q \rangle$ in a way that fixes each ℓ_i, so they must have the same set of values in each block. (This would not be the case if we considered their norm in a global sense, and the restriction modulo q is strictly necessary.) □

Note that the statement does not define the resulting parameter of the Gaussian outputting ℓ, but the proof allows one to compute this: say each k_i was chosen from a discrete Gaussian of parameter r. Then each element of ℓ has parameter $\sqrt{\sum_i |a_i|^2} \cdot r$. Computing $\sqrt{\sum_i |a_i|^2}$ is a one time cost for a fixed choice of $\ell_1, \ldots, \ell_d$, so one can sample the required Gaussian over $\mathcal{O}_{Lq}$ of parameter r' by sampling from the discrete Gaussian over $\mathcal{O}_{Kq}$ of parameter $r = r'/\sqrt{\sum_i |a_i|^2}$.

Finally, to sample elements of Λ_q we merely sample each u coordinate independently according to the above technique. If we wanted to use this method in the cryptosystem of Sect. 5.2 to attain efficient operations then we would sample and store all elements using this representation over the cyclic basis $\ell_1, \ldots \ell_d$.

Unfortunately, we are unable to generalize this theorem to the case where $\mathfrak{q}_i$ remains prime, or even intermediate cases. In this case, there exist cyclic bases of $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ over $\mathcal{O}_K/\mathfrak{q}_i$, but since $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ is a finite field and thus has no zero-divisors the cyclic bases are not orthogonal. Consequentially, the matrix $\mathbf{A}$ does not in general give a diagonal $\mathbf{AA}^\dagger$ and thus the distribution of $\mathbf{Ak}$ has several potentially large covariance terms. If one were able to tolerate the covariance, the method can be extended in this case. It is also possible that a cyclic basis satisfying the condition that $\mathbf{AA}^\dagger$ is diagonal may exist for certain choices of field, but we were not able to find such a family of fields. We note that this question can be asked as a more generic question about finite fields: let $F = \mathbb{F}_{q^d}$ be a finite field with $d > 1$ and let θ denote the Frobenius automorphism of F. Does there exist a cyclic basis $b_1, \ldots, b_d$ with $b_j = \theta^j(b_1)$ for F over $\mathbb{F}_q$ satisfying

$$\sum_{i=0}^{d-1} \theta^i(b_1 \cdot \theta^{j-k}(b_1)) = 0$$

for all $j \neq k$ less than d? Here j and k correspond to j, kth entry of $\mathbf{A}\mathbf{A}^{\dagger}$. We were unable to come up with a basis satisfying this condition, but neither can we show that no such basis exists.

Example 4. We exhibit an example of the basis ℓ_1, ℓ_2 in the simplest setting, that of a degree 2 extension of $\mathbb{Q}$. Let $L = \mathbb{Q}(i)$, with ring of integers $\mathcal{O}_L = \mathbb{Z}[i]$, and consider the ideal $\langle 5 \rangle$ of $\mathcal{O}_L$. 5 factorizes in $\mathcal{O}_L$ as $5 = (2+i)(2-i)$, and it is clear that $\langle 5 \rangle = \langle 2+i \rangle \cdot \langle 2-i \rangle$ is a decomposition into a product of prime ideals.

Using the notation $\mathfrak{q}_1 := \langle 2+i \rangle$, $\mathfrak{q}_2 := \langle 2-i \rangle$, it is easy to check that $2+i = -1 \mod \mathfrak{q}_2$ and thus $-(2+i) = -2-i$ is a valid choice for ℓ_1. Similarly, $-(2-i) = -2+i$ is an appropriate choice for ℓ_2. Correspondingly, the distribution obtained by sampling $k_1, k_2 \leftarrow D_r$, the discrete Gaussian of parameter r over $\mathbb{Z}_5$, and outputting $k_1 \cdot (-2+i) + k_2 \cdot (-2-i)$ is a discrete Gaussian over $\mathcal{O}_L \mod \langle 5 \rangle$. Furthermore, to multiply two elements $k = k_1\ell_1 + k_2\ell_2$ and $g = g_1\ell_1 + g_2\ell_2$ modulo 5 one outputs $kg = (k_1 g_1 \mod 5) \cdot \ell_1 + (k_2 g_2 \mod 5) \cdot \ell_2$, at a cost of two operations in $\mathbb{Z}_5$, and performing the $\mathcal{O}_L \mod 5$ CRT on each u coordinate of an element of the resulting natural order Λ_5 can be done by merely reading off the $d^2 = 4$ values of k_i and no additional computation.

Furthermore, this is an example where the techniques of our next section may be advantageous. We will generalize the multiplication and CRT technique so that one is free to use any basis of $\mathcal{O}_L$ over $\mathbb{Z}$, for example the basis $\{1, i\}$. In this basis it is particularly easy to sample a discrete Gaussian in the polynomial representation of $\mathcal{O}_L \mod \langle 5 \rangle \cong \frac{\mathbb{Z}_5[x]}{x^2+1}$, but the resulting multiplication operation and CRT decomposition is not coordinatewise in the basis and so a small amount of efficiency is lost at a gain in parameter of the Gaussian. Specifically, to compute the CRT on an element $k = k_1 + k_2 \cdot i$, one has to precompute[12] the values $i = -2 \mod \mathfrak{q}_1$, $i = 2 \mod \mathfrak{q}_2$ and output

$$(k_1 - 2k_2 \mod \mathfrak{q}_1, 2k_2 \mod \mathfrak{q}_2),$$

which requires additional operations over $\mathbb{Z}_5$.

F.5. *Generalizing to non-Split q and Arbitrary Bases*

In order to construct the cyclic, orthonormal, basis of Theorem 11, the previous section requires that q be completely split in both K and L. However, it is possible to drop the splitting condition in L and obtain fast multiplication algorithms in the general case at only a small loss of efficiency. We demonstrate the technique in this section and then briefly describe cases where a general algorithm may be superior to the one requiring that q splits by discussing alternatives to Theorem 11.

Observe that, regardless of the prime ideal decomposition of each $\mathfrak{q}_i\mathcal{O}_L$, under the CRT decomposition the quotient ring $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ is a vector space of dimension d over $\mathbb{F}_q \cong \mathcal{O}_K/\mathfrak{q}_i$. Consequentially, an arbitrary $\mathcal{O}_{Kq}$ basis $\ell_1, \ldots, \ell_d$ of $\mathcal{O}_{Lq}$ can be decomposed into n bases $\ell_j = (\ell_{1,j}, \ldots, \ell_{n,j})$ so that each collection $\ell_{i,1}, \ldots, \ell_{i,d}$ of $\mathfrak{q}_i\mathcal{O}_L$ parts is a vector space basis of dimension d over $\mathcal{O}_K/\mathfrak{q}_i$. Indeed, in the split

[12] Note that precomputing the image of 1 is trivial.

case we constructed each ℓ_i in this manner. Armed with this knowledge, we adapt the multiplication algorithm as follows.

Choose an arbitrary integral $\mathcal{O}_K$-basis $\ell_1, \ldots, \ell_d$ of $\mathcal{O}_L$. As a precomputation phase, compute and store the images $\ell_j \mod \mathfrak{q}_i\mathcal{O}_L$ for each i and j. The CRT-like decomposition of Lemma 12 splits each of the u coordinates of an element of Λ_q, an element of $\mathcal{O}_{Lq}$, into its mod $\mathfrak{q}_i\mathcal{O}_L$ parts. Once again, we suggest an algorithm where elements of $\mathcal{O}_{Lq}$ are stored in the form $\ell = \sum_{j=1}^{d} \ell_j k_j$ for $k_j \in \mathcal{O}_{Kq}$, e.g., on elements stored as K-combinations of this basis. We split $\ell \in \mathcal{O}_{Lq}$ into its $\mathcal{O}_L/\mathfrak{q}_i$ components in time $O(d \cdot n \log n)$, since

$$\sum_{j=1}^{d} \ell_j k_j \mod \mathfrak{q}_i\mathcal{O}_L = \sum_{j=1}^{d} (\ell_j \mod \mathfrak{q}_i\mathcal{O}_L) \cdot (k_j \mod \mathfrak{q}_i\mathcal{O}_L),$$

where each $k_j \mod \mathfrak{q}_i$ can be computed in time $O(n \log n)$ by the K-CRT and each $\ell_j \mod \mathfrak{q}_i \mod \mathcal{O}_L$ was computed in the precomputation phase. Consequentially, we can perform the CRT style decomposition of an element in Λ_q whose u coordinates are all stored in this manner in time $O(d^2 \cdot n \log n)$, since we must split d^2 elements of $\mathcal{O}_K$. This decomposing complexity is the same as in the previous case where q splits completely. Following this, each ring $\mathcal{R}_i$ can be plugged in to the algorithm of [15] to compute the multiplication in time $\tilde{O}(Nd^{\omega-2})$. However, since the ℓ_i do not correspond to a standard orthonormal basis we incur an extra cost when reversing this transformation. Namely, each of the u coordinates of each ring $\mathcal{R}_i$ is output by the algorithm of [15] as an element $\ell \in \mathcal{O}_L \mod \mathfrak{q}_i\mathcal{O}_L$ expressed in an arbitrary normal basis. Before reversing the decomposition we must allow for the complexity of expressing each element of the output in the bases obtained by the images of $\ell_1, \ldots, \ell_d \mod \mathfrak{q}_i\mathcal{O}_L$, as this basis was not necessarily normal. Since $\mathcal{O}_L \mod \mathfrak{q}_i\mathcal{O}_L$ is a vector space of dimension d over $\mathbb{F}_q$ this can be done via a precomputed change of basis matrix over $\mathbb{F}_q$ in time $\tilde{O}(d^\omega)$, and since there are n rings with d coordinates each the complexity of computing this on every coordinate is $\tilde{O}(nd^{\omega+1})$. The resulting multiplication algorithm has total complexity $O(N \log(N/d^2)) + \tilde{O}(Nd^{\omega-1})$. While this represents only a minor asymptotic loss, especially since we expect the first term to dominate the complexity, it is likely in practice that the extra step required to recover the basis representation would cause a tangible slowdown.

An unfortunate issue with this technique is that by replacing the orthonormal basis with an arbitrary basis we have lost Theorem 11 and thus the efficient method for sampling a discrete Gaussian in the representation $\ell = \sum_j \ell_j k_j$. However, this generalization allows for the use of an arbitrary basis $\ell_1, \ldots, \ell_d$, unlike in the split case in which we chose a specific basis. Since we require that elements of Λ_q are input into the algorithm with u coordinates in the form $\sum_j \ell_j k_j$ this algorithm can be combined with the cryptosystem of Sect. 5.2 in the case where there is a basis $g_1, \ldots, g_d$ of $\mathcal{O}_{Lq}$ over $\mathcal{O}_{Kq}$ in which one can compute the representation $\ell = \sum_j g_j k_j$ particularly efficiently. This is because one can just sample ℓ from the usual Gaussian distribution over the polynomial basis of $\mathcal{O}_{Lq}$, compute its representation as $\ell = \sum_j g_j k_j$, and then apply the multiplication algorithm in this form. More generally, the flexible choice of basis allows for both non-split q and for a user to choose their favorite $\mathcal{O}_L$ basis properties,

such as a normal basis or a basis consisting of small elements. We remark that it is likely possible to construct a pair of fields L/K that allow for a basis $\ell_1, \ldots, \ell_d$ permitting a fast algorithm transforming from the polynomial representation of $\mathcal{O}_L$ to the representation $\sum_i \ell_i k_i$ with each k_i in polynomial representation, which would allow one to bypass the complications of sampling Gaussian distributions by just sampling in $\mathcal{O}_L$ directly.

F.6. *Generalizing to Other Centers*

In the exposition of the previous section we required that q splits completely in the center K. This corresponds to the requirement in the ring and module cases that q splits completely in the field K, which allows the use of the NTT to compute multiplications over a direct product of finite fields. However, there has been recent progress in loosening this requirement for the NTT and allowing the modulus q to be 1 mod n rather than 1 mod m, where as usual K is the mth cyclotomic field of degree n. For example, in the second round specification of KYBER [5] q is set as 3329 and $n = 256$, yet they still support efficient NTT-based multiplication. In such cases, q is 'well' split but not completely split, and the fast NTT operations use the method of [29], where q splits into some product of prime ideals $\mathfrak{q}_i$ whose norms can be small powers of q.

We observe that our methods can be partially generalized to this case in the following manner. Say $\langle q \rangle = \prod_i \mathfrak{q}_i$ is a decomposition into prime ideals in $\mathcal{O}_K$ and there exists an efficient algorithm for fast multiplication in $\mathcal{O}_{Kq}$. We can replace our condition that q splits completely in $\mathcal{O}_L$ with the condition that each ideal $\mathfrak{q}_i$ in the $\mathcal{O}_K$-factorization of q splits completely into a product of d prime ideals $\mathfrak{q}_i \mathcal{O}_L = \prod_{j=1}^{d} \mathfrak{q}_{i,j}$ in $\mathcal{O}_L$ of the same norm. Then, we can replicate the method of Appendix F.4 to find a cyclic, orthonormal basis $\mathbf{e}_1, \ldots, \mathbf{e}_d$ of $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ over $\mathcal{O}_K/\mathfrak{q}_i$ and concatenate together the bases for each i to make the cyclic, orthonormal, basis $\ell_1, \ldots, \ell_d$ of $\mathcal{O}_{Lq}$ over $\mathcal{O}_{Kq}$. Since the basis is orthonormal, if $\ell = \sum_i \ell_i k_i$ and $g = \sum_i \ell_i g_i$ with each $k_i, g_i \in \mathcal{O}_{Kq}$, then

$$\ell \cdot g = \sum_{i=1}^{d} \ell_i (g_i \cdot k_i).$$

Since the basis is cyclic,

$$\begin{aligned} \theta(\ell) &= \sum_i \theta(\ell_i) k_i \\ &= \sum_i \ell_i k_{i-1} \end{aligned}$$

where we define $k_0 := k_d$.

Now we are able to use existing fast multiplication algorithms in $\mathcal{O}_{Kq}$ to compute operations in $\mathcal{O}_{Lq}$ by expressing elements in this basis. Represent each $x = \sum_{i=0}^{d-1} u^i x_i \in \Lambda_q$ by expressing each $x_i \in \mathcal{O}_{Lq}$ in the ℓ_j basis. Then, to multiply x and y in Λ_q one only has to compute multiplications in $\mathcal{O}_{Kq}$, since the operations required are just computing the non-commutative relation $\ell u = u\theta(\ell)$, which merely permutes the ℓ_i using θ,

and computing multiplication and addition, which can be done coordinatewise in the orthonormal ℓ_i basis. Each L multiplication requires d multiplications in K, and each u coordinate of Λ requires d multiplications in L. Consequentially, naive multiplication in Λ_q takes d^3 instances of the efficient $\mathcal{O}_{Kq}$-multiplication algorithm we have access to. For specific K-multiplication algorithms it is likely that this process can be streamlined; the intention of this section is merely to demonstrate that one can build efficient Λ_q operations from more general efficient operations over the center in the same manner that the techniques of Appendix F.4 used the CRT method.

References

[1] G. Alagic, J. Alperin-Sheriff, D. Apon, D. Cooper, Q. Dang, J. Kelsey, Y.K. Liu, C. Miller, D. Moody, R. Peralta, R. Perlner, A. Robinson, D. Smith-Tone, Status report on the second round of the NIST post-quantum cryptography standardization process. Tech. rep., NIST (2020), https://nvlpubs.nist.gov/nistpubs/ir/2020/NIST.IR.8309.pdf

[2] M.R. Albrecht, A. Deo, Large modulus Ring-LWE ≥ Module-LWE, in Takagi, T., Peyrin, T. (eds.) *Advances in Cryptology – ASIACRYPT 2017*. pp. 267–296. Springer, Cham (2017)

[3] E. Alkim, L. Ducas, T. Pöppelmann, P. Schwabe, Post-quantum key exchange: a new hope, in *25th USENIX Security Symposium (USENIX Security 16)*. pp. 327–343 (2016)

[4] B. Applebaum, D. Cash, C. Peikert, A. Sahai, Fast cryptographic primitives and circular-secure encryption based on hard learning problems, in *Advances in Cryptology-CRYPTO 2009*, pp. 595–618. Springer (2009)

[5] R. Avanzi, J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J.M. Schanck, P. Schwabe, G. Seiler, D. Stehlé, CRYSTALS-Kyber algorithm specifications and supporting documentation (version 2.0). https://pq-crystals.org/kyber/data/kyber-specification-round2.pdf (2019)

[6] W. Banaszczyk, New bounds in some transference theorems in the geometry of numbers. *Math. Annalen* **296**(1), 625–635 (1993)

[7] A. Banerjee, C. Peikert, New and improved key-homomorphic pseudorandom functions. in Garay, J.A., Gennaro, R. (eds.) *Advances in Cryptology – CRYPTO 2014*. pp. 353–370. Springer, Berlin, Heidelberg (2014)

[8] G. Baumslag, N. Fazio, A.R. Nicolosi, V. Shpilrain, W.E. Skeith III, Generalized learning problems and applications to non-commutative cryptography, in *Provable Security*, pp. 324–339. Springer (2011)

[9] G. Berhuy, F. Oggier, *An Introduction to Central Simple Algebras and Their Applications to Wireless Communication*. American Mathematical Society (2013)

[10] J.F. Biasse, F. Song, On the quantum attacks against schemes relying on the hardness of finding a short generator of an ideal in Q (ζ_p^n). Tech. rep. (2015)

[11] M. Bolboceanu, Z. Brakerski, R. Perlman, D. Sharma, Order–LWE and the hardness of Ring–LWE with entropic secrets. Cryptology ePrint Archive, Report 2018/494 (2018), https://eprint.iacr.org/2018/494

[12] C. Bootland, W. Castryck, F. Vercauteren, On the Security of the Multivariate Ring Learning with Errors Problem (2018), published: Cryptology ePrint Archive, Report 2018/966

[13] J. Bos, C. Costello, L. Ducas, I. Mironov, M. Naehrig, V. Nikolaenko, A. Raghunathan, D. Stebila, Frodo: Take off the ring! Practical, Quantum-Secure Key Exchange from LWE (2016), published: Cryptology ePrint Archive, Report 2016/659

[14] P. Campbell, M. Groves, D. Shepherd, Soliloquy: A cautionary tale (2015)

[15] X. Caruso, J. Le Borgne, Fast multiplication for skew polynomials, in *Proceedings of the 2017 ACM on International Symposium on Symbolic and Algebraic Computation*, pp. 77–84. ACM (2017)

[16] Q. Cheng, J. Zhuang, LWE from Non-commutative Group Rings. arXiv preprint arXiv:1612.06670 (2016)

[17] R. Cramer, L. Ducas, C. Peikert, O. Regev, Recovering short generators of principal ideals in cyclotomic rings, in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. pp. 559–585. Springer (2016)

[18] R. Cramer, L. Ducas, B. Wesolowski, Short Stickelberger class relations and application to Ideal-SVP. in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. pp. 324–348. Springer (2017)

[19] E. Crockett, C. Peikert, Challenges for Ring-LWE. IACR Cryptology ePrint Archive (2016)

[20] R. Jozsa, Quantum factoring, discrete logarithms, and the hidden subgroup problem. *Comput. Sci. Eng.* **3**(2), 34–43 (2001)

[21] J. Lahtonen, N. Markin, G. McGuire, Construction of multiblock space–time codes from division algebras with roots of unity as nonnorm elements. *IEEE Trans. Inf. Theory* **54**(11), 5231–5235 (2008)

[22] A. Langlois, D. Stehlé, Worst-case to average-case reductions for module lattices. Designs Codes Cryptogr. **75**(3), 565–599 (2015)

[23] H. Lu, Constructions of multiblock space–time coding schemes that achieve the diversity-multiplexing tradeoff. IEEE Trans. Inf. Theory **54**(8), 3790–3796 (2008)

[24] X. Lu, X. Liu, Z. Zhang, D. Jia, H. Xue, J. He, B. Li, K. Wang, Z. Liu, H. Yang, LAC: Practical Ring–LWE based public-key encryption with byte-level modulus (2018), https://eprint.iacr.org/2018/1009.pdf

[25] L. Luzzi, R. Vehkalahti, C. Ling, Almost universal codes for MIMO wiretap channels. IEEE Trans. Inf. Theory **64**(11), 7218–7241 (2018)

[26] V. Lyubashevsky, C. Peikert, O. Regev, On ideal lattices and learning with errors over rings, in Gilbert, H. (Ed.) *Advances in Cryptology – EUROCRYPT 2010*, (Springer, Berlin, Heidelberg, 2010), pp. 1–23.

[27] V. Lyubashevsky, C. Peikert, O. Regev, On ideal lattices and learning with errors over rings. in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, (Springer, 2010), pp. 1–23.

[28] V. Lyubashevsky, C. Peikert, O. Regev, A toolkit for Ring-LWE cryptography. in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, (Springer, 2013), pp. 35–54.

[29] V. Lyubashevsky, G. Seiler, NTTRU: truly fast NTRU using NTT. *IACR Trans. Cryptogr. Hardware Embed. Syst.* **2019**(3), 180–201 (2019)

[30] C. Maire, F. Oggier, Maximal order codes over number fields. J. Pure Appl. Algebra **222**(7), 1827 – 1858 (2018)

[31] D. Micciancio, O. Regev, Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.* **37**(1), 267–302 (2007)

[32] F. Oggier, J.C. Belfiore, E. Viterbo, *Cyclic Division Algebras: A Tool for Space-time Coding*, (Now Publishers Inc, 2007)

[33] F. Oggier, B.A. Sethuraman, Quotients of orders in cyclic algebras and space-time codes. *Adv. Math. Commun.*, **7** (2012)

[34] F. Oggier, B. Sethuraman, *Quotients of orders in cyclic algebras and space-time codes*. arXiv preprint arXiv:1210.7044 (2012)

[35] A. Pedrouzo-Ulloa, J.R. Troncoso-Pastoriza, F. Pérez-González, On Ring Learning with Errors over the Tensor Product of Number Fields. arXiv preprint arXiv:1607.05244 (2016)

[36] A. Pedrouzo-Ulloa, J.R. Troncoso-Pastoriza, N. Gama, M. Georgieva, F. Pérez-González, Revisiting multivariate ring learning with errors and its applications on lattice-based cryptography. Cryptology ePrint Archive, Report 2019/1109 (2019), https://eprint.iacr.org/2019/1109

[37] C. Peikert, An efficient and parallel Gaussian sampler for lattices. in: *Annual Cryptology Conference*, (Springer, 2010), pp. 80–97

[38] C. Peikert, How (not) to instantiate ring-LWE. in *International Conference on Security and Cryptography for Networks*, (Springer, 2016), pp. 411–430

[39] C. Peikert, Z. Pepin, Algebraically structured LWE, revisited. Cryptology ePrint Archive, Report 2019/878 (2019), https://eprint.iacr.org/2019/878

[40] C. Peikert, O. Regev, Stephens-Davidowitz, N.: Pseudorandomness of ring-LWE for any ring and modulus. in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, (ACM, 2017), pp. 461–473

[41] R.S. Pierce, *Associative algebras*. Graduate Texts in Mathematics, (Springer, New York, NY 1982)

[42] O. Regev, On lattices, learning with errors, random linear codes, and cryptography. *J. ACM (JACM)* **56**(6), 34 (2009)

[43] I. Reiner, *Maximal Orders. L.M.S. Monographs*. Academic Press (1975)

[44] R. Vehkalahti, C. Hollanti, J. Lahtonen, K. Ranto, On the densest MIMO lattices from cyclic division algebras. *IEEE Trans. Inf. Theory* **55**(8), 3751–3780 (2009)
[45] L.C. Washington, *Introduction to Cyclotomic Fields. Graduate Texts in Mathematics*, (Springer, New York, 2012)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.