Structured Module Learning With Errors From Cyclic Algebras

Charles Grover, Cong Ling

Imperial College London
c.grover15@imperial.ac.uk, c.ling@imperial.ac.uk

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 LWE instance. 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. The proposed construction is both more efficient than Module LWE and conjecturally more secure than Ring LWE, the best of both worlds. We show that the standard 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. As a contribution of theoretic interest, we view CLWE as the first variant of LWE which naturally supports non-commutative multiplication operations.

1 Introduction

With the predicted advent of quantum computers compromising the bulk of existent cryptographic constructions, lattice based cryptography has emerged in the last ten years as a promising foundation for long term security. In particular, the Learning with Errors (henceforth LWE) problem introduced in [26], as well as its variants over rings (RLWE) [16] and modules (MLWE) [14], 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 standardisation process such as NewHope [1] and KYBER [6] 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.

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 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 $A, \mathbf{b}$. The ring variant replaces $A, \mathbf{s}, \mathbf{e}$ with elements a, s, e from the ring $R_q := \dfrac{\mathbb{Z}_q[x]}{x^n + 1}$, 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 $a = a_0 + a_1 x + ... + a_{n-1} x^{n-1}$ and bold faced letters denoting coefficient vectors, a sample from the RLWE distribution takes the form:

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

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 $\mathbf{A}$ has been replaced by the negacyclic matrix $\mathrm{rot}(a)$. Of course, this should be an easier problem 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 $\mathrm{rot}(\cdot)$ 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 $\mathrm{rot}(a)$ 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.

This concept of saving memory by adding structure motivates this work; can we perform an analog of the transformation taking an LWE matrix A to an RLWE matrix $\mathrm{rot}(a)$ 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. [15, 18, 28]) 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 space. Though some efforts have been made to construct non-commutative LWE problems, for example [3], [9], 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 $\mathrm{rot}(\cdot)$ operation coordinate-wise, 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 ideals in a cyclic algebra to CLWE, establishing its security on the assumption that such problems are hard. Specifically, due to their similarity with the concept of ideals in the ring of integers of a number field K, we consider ideal lattices induced by the so-called natural order of $\mathcal{A}$.

The 'standard' security reductions used in [26] and [16] use similar machinery to reduce search LWE and RLWE to their respective lattice problems, then establish hardness of the decision problem (the problem of distinguishing LWE samples $a \cdot s + e$ from the uniform distribution) via a search-decision reduction. We reduce search CLWE to a BDD problem using the same method as in [16]. 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 Chinese Remainder Theorem. Though this is not immediately suitable for our needs, because the relative factorization of ideals required does not transfer to ideals in the natural order, we apply a similar method in suitable subrings to deduce the hardness of our decision problem. As in [16], the algorithm bases the security of CLWE on short vector problems over 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.

We consider the advantages of our CLWE construction:

- Efficiency. 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. Multiplication in cyclic algebras can be implemented over a product of skew polynomial rings following a CRT style decomposition, for which well known fast algorithms, such as those of [8] and [25], can be combined with the decomposition of our Lemma 13 to compute the operation $A \cdot \mathbf{s}$ more efficiently in the case where $A = \phi(a)$ than in the module case where A is uniform.
- Security. Following recent works on quantum attacks on related ideal lattice problems (e.g. [4], [10], [11], [7] amongst others), we observe that the non-commutativity of multiplication in cyclic algebras may be viewed as a security advantage. This is because the Hidden Subgroup Problem (HSP), an integral part of the majority of algorithms using quantum computing to gain an advantage over classical computation, requires that the underlying group, in this case the unit group of $\mathcal{O}_K$, is commutative, see e.g. [12], which is untrue for a non-commutative algebra. We conjecture that the security level will be as high as MLWE, but welcome further cryptanalysis. We actively avoid known attacks on previous attempts to create structured MLWE (see Appendix A).
- Functionality. The non-commutative algebraic structure of CLWE opens up the prospect of extra functionality. For example, since operations are composable and non-commutative, one could hope to construct FHE in a non-commutative space. We leave this frontier open for separate work.

Overall we consider it plausible that LWE in cyclic algebras could be both more efficient than MLWE and more secure than Ring LWE in a quantum setting

1.2 Related Work and Future Work

An alternative construction for structured module LWE, called multivariate-RLWE, was presented in [20], where they tensor product two (or more) number fields in order to provide a structured module matrix. However, their efficient implementations were attacked in [5], together with a warning about taking care

when putting structure on a module. In short, [5] attacks certain instances of m-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 Appendix A).

As for 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, we were unable to show a direct-to-decision reduction by adapting the methods of [23], which may generalize the choice of algebras. Finally, this work is focused on the theoretical construction of a non-commutative LWE assumption, and we leave practical analysis and implementation of cryptography based on CLWE as further research.

Roadmap In Section 2 we provide necessary background material on lattices, number fields, and cyclic algebras. In Section 3 we provide a definition and discussion of Cyclic LWE. In Section 4 we provide a reduction from search CLWE to structured lattice problems. In Section 5 we provide search-worst case decision reduction for CLWE. Finally, in Section 6 we show a normal form reduction for CLWE and provide a sample cryptosystem.

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 $B = \{\mathbf{b}_1, ..., \mathbf{b}_k\}$ 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 $\lambda_1(\mathcal{L}) = \min_{\boldsymbol{v} \in \Lambda/\{0\}} \|\boldsymbol{v}\|$. 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 $\{\boldsymbol{x}_1, ..., \boldsymbol{x}_n\}$ is defined as $\max_i(\|\boldsymbol{x}_i\|)$.*

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 $\mathcal{L}^* = \{\boldsymbol{v} \in V : \langle \mathcal{L}, \boldsymbol{v} \rangle \subset \mathbb{Z}\}$.*

We will return to lattices later, once we have defined some particular vector spaces of interest.

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 $\rho_r(\boldsymbol{x}) = \exp(-\pi\|\boldsymbol{x}\|/r^2)$.*

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 $\mathbf{b}_1, ..., \mathbf{b}_n$ of V, for $\mathbf{r} = (r_1, ..., r_n)$ a vector of positive reals, by sampling $x_1, ..., x_n$ 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 $\dfrac{\rho_r(\mathbf{x})}{\rho_r(\mathcal{L})}$ 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 [17]. The motivation for the name is provided by Lemma 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 $\rho_{1/r}(\mathcal{L}^*/\{\boldsymbol{0}\}) \leq \varepsilon$.*

The following is a special case of [17], 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, $\rho_r(\mathcal{L} + \boldsymbol{x}) \in [\frac{1-\varepsilon}{1+\varepsilon}] \cdot \rho_r(\mathcal{L})$.*

We introduce well known lemmas used to relate the smoothing parameter to standard lattice properties. The first comes from [2], the second from [23].

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 $\eta_\varepsilon(\mathcal{L}) \geq \dfrac{\sqrt{\log(1/\varepsilon)/\pi}}{\lambda_1(\mathcal{L}^*)}$.*

2.3 Algebraic Number Theory

Basic Definitions

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 n. The Galois group of L over K is the group $Aut(L/K)$ of automorphisms of L that fix K. We say that the extension is Galois if the subfield of L fixed by $Aut(L/K)$ 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 $n := [L : K]$. 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 $\mathcal{O}_L \cap K = \mathcal{O}_K$.

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*

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

Formally, σ maps into the space

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

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.

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 k and L has degree n over K, there are precisely k embeddings $\sigma_1, ..., \sigma_k$ of K into $\mathbb{C}$. Assuming $L/\mathbb{Q}$ is Galois, each of these can be extended to an embedding $\alpha_i : L \to L$ such that $\alpha_i|_K = \sigma_i$. However, these extensions are not unique, and it is easy to see that there are $[L : K] = n$ 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 $\alpha_i \circ \theta^j(\cdot), 1 \leq j \leq n$, run through the n choices of α_i. Hence for a fixed choice of $\alpha_1, ..., \alpha_k$ the nk automorphisms of L can each be uniquely represented by some $\alpha_i \circ \theta^j(\cdot)$, which we denote by $\alpha_i^j(\cdot), 1 \leq i \leq k, 1 \leq j \leq n$. 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 k, and let L be a Galois extension of K of degree n such that the Galois group of L over K is cyclic of degree n, $Gal(L/K) = \langle\theta\rangle$. For non-zero $\gamma \in K$ we define the resulting cyclic algebra*

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

where $u \in \mathcal{A}$ is some auxiliary generating element of $\mathcal{A}$ satisfying the additional relations $xu = u\theta(x)\, \forall x \in L$ and $u^n = \gamma$. We will call n 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 $aa^{-1} = 1$.

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.

Matrix Representation We present a 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 n-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 $\phi : \mathcal{A} \to M_{n\times n}(L)$, where for $x = x_0 + ux_1 + ... + u^{n-1}x_{n-1} \in$

$\mathcal{A}$ with each $x_i \in L$,

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

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 $\phi : \bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}} \to M_{n\times n}(L_{\mathbb{R}})$. We derive lattices from subrings of a cyclic algebra by vectorising their images under ϕ.

Orders and the Natural Order

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 $\Lambda := \bigoplus_{i=0}^{n-1} u^i \mathcal{O}_L$. Unlike in the case of $\mathcal{O}_K$, this order is not necessarily maximal. Note that in order for Λ to be closed under multiplication the element γ must lie in $\mathcal{O}_K$.

Existence and Construction It is not a priori obvious whether well-defined cyclic 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', which may be found in [28]. This condition states that the lowest power of γ that appears in $N_{L/K}(L)$, is γ^n, where $N_{L/K}$ represents the relative norm of L into K.

Order Ideals Analogous to the use of $\mathcal{O}_K$ ideals in RLWE, we will be interested in ideals of the natural 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 $c \in K$.

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, Theorem 22.10 of [27] states that the group of two-sided ideals is a free abelian group generated by the prime (e.g. maximal) ideals, from which one can deduce obvious definitions of inverse and coprime ideals. For a general order Γ, we define its prime ideals as its maximal two-sided ideals and the inverse of an ideal $\mathcal{I} \subset \Gamma$ is

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

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, and so rely on the exposition of Section 3 of [19]. In particular, we will use the fact that 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 $\mathcal{A}$ 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 $\mathcal{I} + \mathcal{J} = \Lambda$.

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 arbitrary two-sided 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.

Orders and Ideals as Integer Lattices Any order Λ of a cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ has dimension kn^2 over $\mathbb{Z}$ and thus generates a lattice of dimension kn^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 $x = \bigoplus_{i=0}^{n-1} u^i x_i \in \Lambda$. We can consider x as a vector over H_L of dimension n by $\sigma_{\mathcal{A}}(x) := \{\sigma_L(x_0), \sigma_L(x_1), ..., \sigma_L(x_{n-1})\}$. Then, the collection $\sigma_{\mathcal{A}}(\Lambda)$ forms an integer lattice of dimension kn^2. 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 $\bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}}$. 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.

We will also make use of the following representation induced by the matrix embedding. First, recall that for any embedding $\sigma_i : K \to K$ there are precisely n extension embeddings $\alpha_i^j : L \to L$ which agree with σ_i on K, and for a fixed choice of α_i the compositions $\theta^j \circ \alpha_i : L \to L, j = 1, ..., n$ run through all such choices. Fixing some choice $\alpha_1, ..., \alpha_k$ of extensions of the embeddings of K and letting $\alpha(y) = (\alpha_1(y), ..., \alpha_k(y))^T$ be a column vector for $y \in L$, we denote by $\alpha(\phi(x))$ the $nk \times n$ matrix obtained by applying α coordinatewise to $\phi(x)$. Stacking the columns of this matrix results in a vector of dimension kn^2; applying this function to all elements of an order gives an integer lattice. Indeed, since $\theta^j \circ \alpha_i(\cdot)$ runs over all of the embeddings of L this is just a reordering of the entries of the previous visualization, with the exception that some elements have been multiplied by various α embeddings of γ. We remark here that since $\alpha_i^j(\cdot)$ agree on K for fixed i it follows that $\alpha_i(\gamma)$ is fixed over the choice of α_i.

Gaussian Distributions Over Cyclic Algebras 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 $K_{\mathbb{R}} \cong H$, defined in Section 2.3. Since a cyclic algebra $\mathcal{A}$ can be viewed as an n-dimensional algebra over L, we use the first visualization from the previous subsection and sample our error distributions over $\bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}}$, which has the same structure as a vector space as ${H_L}^n$. 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. Otherwise, in order to maintain a norm that is sub-multiplicative the norm and shape of γ must be considered.

Explicitly, we just consider the norm of an element of $\mathcal{A}$ to be equal to the norm of the corresponding module element in L^n of dimension kn^2 used in [14], e.g. $\|x\| = \|(\sigma_L(x_0), \sigma_L(x_1), ..., \sigma_L(x_{n-1}))\|_2$ for $x = x_0 + ux_1 + ... + u^{n-1}x_{n-1} \in \mathcal{A}$. 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. It is clear that this norm extends to any $y \in \bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}}$ 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 ${L_{\mathbb{R}}}^n$.

The Chinese Remainder Theorem In this subsection we state the Chinese Remainder Theorem for order ideals, and deduce some important consequences. We note that the following lemmas are merely adaptations of those in Section 2.3.8 of [16] extended to the case of cyclic algebras. The first is just the Chinese Remainder Theorem.

Lemma 4. *Let $\mathcal{I}_1, ..., \mathcal{I}_r$ be pairwise coprime two-sided 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 $\mathcal{I}_1, ..., \mathcal{I}_r$ a basis $C = \{c_1, ..., c_r\}$ 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 two-sided ideals $\mathcal{I}_1, ..., \mathcal{I}_r$ of a maximal order Λ, there is a deterministic polynomial time algorithm that outputs a CRT basis $c_1, ..., c_r \in \Lambda$ for those ideals.*

The proof is the same as in the ring case, Lemma 2.13 of [16]. Using Lemma 5 we can efficiently invert the natural CRT isomorphism. Given $a = (a_1, ..., a_r) \in \bigoplus_{i=1}^r (\Lambda/\mathcal{I}_i)$, 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. *Let $\mathcal{I}, \mathcal{J}$ be two-sided ideals of the natural order Λ. 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 $\mathcal{J} \cap \mathcal{O}_K$.*

Proof. For an ideal $\mathcal{I} \subset \Lambda$, denote by $\overline{\mathcal{I}}$ its intersection with $\mathcal{O}_K$. We apply the corresponding Lemma 2.14 of [16] to obtain $t \in \overline{\mathcal{I}}$ such that $t \cdot \overline{\mathcal{I}}^{-1}$ and $\overline{\mathcal{J}}$ are coprime as ideals of $\mathcal{O}_K$. Since $\overline{\mathcal{I}}^{-1} \subset \overline{\mathcal{I}^{-1}}$ and $t \cdot \overline{\mathcal{I}}^{-1} + \overline{\mathcal{J}} = \mathcal{O}_K$, we have $t \cdot \overline{\mathcal{I}^{-1}} + \overline{\mathcal{J}} = \mathcal{O}_K$. Since t lies in the center of $\mathcal{A}$, we have $\overline{t \cdot \mathcal{I}^{-1}} = t \cdot \overline{\mathcal{I}^{-1}}$. Now observing that $\overline{t \cdot \mathcal{I}^{-1}} + \overline{\mathcal{J}} \subset \overline{t \cdot \mathcal{I}^{-1} + J}$ we see that $\overline{t \cdot \mathcal{I}^{-1} + J} = \mathcal{O}_K$, from which it follows that $t \cdot \mathcal{I}^{-1} + J = \Lambda$, since the lift of any $\mathcal{O}_K$ ideal $\mathcal{P}$ must contain the ideal $\mathcal{P} \cdot \Lambda$. □

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 $\mathcal{I}, \mathcal{J}$ be two-sided 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 two-sided 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 [16]. 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 $\mathcal{I}^{-1} \cdot t \cdot x \cdot \mathcal{P}^{-1} \subset \mathcal{J}$. 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 $(a + b) \cdot (x \cdot \mathcal{P}^{-1}) = x \cdot \mathcal{P}^{-1}$, and hence that $x \cdot \mathcal{P}^{-1} \subset \mathcal{J}$; finally $x \in \mathcal{J}\cdot\mathcal{P}$, 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 $k^2 \cdot n^4$ elements $x_i, i = 1, ..., k^2 \cdot n^4$ 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 $y'_1, ..., y'_{k\cdot n^2}$. This basis satisfies the desired property that each element has a known inverse. If this algorithm fails (e.g. there is no suitable basis $y'_1, ...y'_{k\cdot n^2}$), we repeat, choosing a fresh set of elements $x_1, ..., x_{k^2\cdot n^4}$ until we succeed. □

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

This lemma is a straightforward adaptation of Corollary 3.16 of [26].

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 non-zero vector $\boldsymbol{x}$ such that $\|\boldsymbol{x}\| \leq \xi \cdot \lambda_1(\mathcal{L})$.*

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

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 ($BDD_{\mathcal{L},d}$) on input $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{e}$ for $\boldsymbol{x} \in \mathcal{L}$ and $\|\boldsymbol{e}\| \leq d$ is to compute $\boldsymbol{x}$, or equivalently $\boldsymbol{e}$.*

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 Cyclic Algebra LWE, 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 $\alpha(\phi(\cdot))$. 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 non-zero element $a \in \mathcal{I}$ such that $|a| := \|\alpha(\phi(x))\|_2 \leq \xi \cdot \lambda_1(\alpha(\phi(\mathcal{I})))$, where as usual $\lambda_1(\mathcal{I})$ denotes the minimal length of elements of $\mathcal{I}$ in the given norm.*

Remark 1. 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. Additionally, in the case where $|\gamma| = 1$ as an element of $\mathbb{C}$, it follows that $\|\alpha(\phi(x))\|_2 = \|\sigma_{\mathcal{A}}(x)\|_2$.

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' $a_1, ..., a_r$ 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 $d < \lambda_1(\mathcal{I})/2$. Then the $\mathcal{A}$-BDD$_{\mathcal{I},d}$ problem, on input $y = x + e$ for $x \in \mathcal{I}$ and $e \in \bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}}$ satisfying $|e| \leq d$, is to compute x.*

2.6 The Learning With Errors Problem

We will briefly recall the initial Learning With Errors (LWE) problem here; in Section 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 $\boldsymbol{s} \in \mathbb{Z}_q^n$, a sample $(\boldsymbol{a}, b) \leftarrow A_{\boldsymbol{s},\alpha}$ is taken by sampling a uniformly random vector $\boldsymbol{a} \in \mathbb{Z}_q^n$ and $e \leftarrow D_\alpha$ and outputting $(\boldsymbol{a}, b) = (\boldsymbol{a}, \langle \boldsymbol{a}, \boldsymbol{s} \rangle / q + e \mod \mathbb{Z})$.*

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

Definition 17. *The search LWE problem is to recover* $\boldsymbol{s}$ *from a collection of samples* $A_{\boldsymbol{s},\alpha}$*. 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* $A_{\boldsymbol{s},\alpha}$ *for some secret* $\boldsymbol{s}$*, providing the samples were taken from one of these distributions.*

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 [16] and [14] respectively, both of which we reference frequently in this work.

LWE Security In his initial LWE paper [26], Regev demonstrated a polynomial time, partly quantum, reduction from SVP (or SIVP) to the Decision LWE problem. Broadly speaking, the algorithm shows that a BDD instance can be converted to an LWE sample. Then, there is a quantum algorithm given to reduce the problem of sampling from a narrow discrete Gaussian distribution to solving a BDD instance. Finally, sampling a narrow discrete Gaussian is sufficient to solve the short vector problem. Thus, solving the search LWE problem is sufficient to solve SVP; it follows that LWE is at least as hard as SVP for suitable parameters. Regev also showed that search LWE and decision LWE are polynomial time equivalent, and provided a worst case-to-average case reduction for these problems. In [16], Lyubashevsky et al. construct analogous reductions from RLWE to lattice problems on ideal lattices in number fields, though their reduction relied on particular choices of number fields and moduli to support the search-to-decision step. [14] does the same for MLWE to module lattice problems.

3 Cyclic Algebra Learning With Errors

In this section we present the general construction of CLWE together with justifications for choices made in the definition, as well as suggestions for specific algebras to use. We will save the security properties for the Section 4.

Definition 18. *Let* L/K *be a Galois extension of number fields of dimension* $[L:K] = n$*,* $[K:\mathbb{Q}] = k$ *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* $u^n = \gamma \in \mathcal{O}_K$*. Let* Λ *be the natural order of* $\mathcal{A}$*. For an error distribution* ψ *over* $\bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}}$*, 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* $(a,b) = (a, (a \cdot s)/q + e \mod \Lambda^\vee) \in \Lambda_q, (\bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}})/\Lambda^\vee$.

Remark 2. Unlike in commutative spaces, the order of multiplication of a and s is important; we will justify our choice momentarily, but it seems likely that similar security properties would hold if one took $(s \cdot a)/q + e$ instead. Also observe that our modulo reduction in the second coordinate of the pair is well defined, since $(a \cdot s)/q \in \Lambda_q^\vee$.

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

Definition 19. *Let $\Pi_{q,s,\psi}$ be a CLWE distribution for parameters $q \geq 2$, $s \in \Lambda_q^\vee$, and error distribution ψ. Then, the search CLWE problem, which we denote by $CLWE_{q,s,\psi}$, is to recover $s \in \Lambda_q^\vee$ from a collection of independent samples from $\Pi_{q,s,\psi}$.*

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 $\bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}}$ and U_Λ denote the uniform distribution on $(\Lambda_q, (\bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}})/\Lambda^\vee)$. Then, the decision CLWE problem, written $CDLWE_{q,\Upsilon}$, 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

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 $a, s, e \in R_q \cong \dfrac{\mathbb{Z}_q[x]}{x^n + 1}$, generates n LWE samples by considering the multiplication operation as $A\mathbf{s} + \mathbf{e}$, where $A := \mathrm{rot}(a)$ 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 $\Lambda_q \cong \bigoplus_{i=0}^{n-1} u^i \mathcal{O}_L / q\mathcal{O}_L$ (see [19]), 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 n 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 n samples from the MLWE distribution with some additional structure in the matrix A, as intended.

The Natural Order We have chosen to use the natural order rather than some maximal order of an algebra $\mathcal{A}$ for a few reasons. Firstly, the natural order 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 $\mathcal{O}_K^d$, and also provides a considerable advantage in terms of conceptual simplicity.

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 k or n, but following our motivation from MLWE we suggest that k is likely the suitable choice since the module dimension n 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 k and n asymptotically; the only case a cyclic algebra precludes is high dimensional MLWE over a low dimension number field L, because the parameter n occurs in both the module and field dimension.

Explicit Examples of Suitable Algebras We demonstrate a method for constructing suitable algebras for CLWE, together with some novel examples of families of algebras we suggest as plausible choices for cryptography. Recall that we seek algebras with root of unity γ in order to simplify our norm.

Lemma 9. *For an integer $m \geq 1$ and odd prime p there exist cyclic division algebras $\mathcal{A} = (L/K, \theta, \gamma)$ such that $K = \mathbb{Q}(\zeta_{2^m \cdot p})$, $[L : K] = p$, and γ is a root of unity.*

Proof. We begin with a construction from Section 4 of [13], which constructs a cyclic division algebra[1] $\mathcal{A} = (F/M, \theta', \gamma)$ where $M = \mathbb{Q}(\zeta_n)$ for some prime power $n := p^a$, γ is a root of unity, and F is a well chosen field such that F/M is a cyclic Galois extension of degree n with $\gamma = \zeta_n$ is a non-norm element. Although this construction provides us with a suitable γ, the relative degrees of the two extensions are close. We proceed to increase the size of the first extension while fixing $[F : M]$. We start by setting $a = 1$.

A consequence of Theorem 30.8 of [27] is that for another extension E/M satisfying $E \cap F = M$, the algebra $\mathcal{A}' = (EF/E, \theta, \gamma)$ is a well-defined cyclic algebra, where θ is any generator of the cyclic Galois group of EF/E of degree $[F : M]$. Furthermore, an easy application of the tower lemma shows that as long as the relative degrees $[E : M]$ and $[F : M]$ are coprime, then $E \cap F = M$, from which it follows that γ is still a non-norm element. Explicitly, combined with the construction of [13] one can proceed with odd prime n e.g. $3, 5, 7$ and field $E := M(\zeta_{2^m}) = \mathbb{Q}(\zeta_{2^m \cdot p})$. Now $[E : \mathbb{Q}] = (n-1) \cdot 2^{m-1}$ and $[EF : E] = [F : M] = n$, and it is clear that the degrees of E and F over M are coprime. The resulting algebra $\mathcal{A}'$ is as required, after relabeling $K := E$ and $L := EF$. □

[1] In fact, they prove the existence of infinitely many such algebras, but we only need one.

In our structured module analogy, this family corresponds to modules of small prime dimension n over number fields of degree $n(n-1)2^{m-1}$ for any choice of integer m. Note that setting $n = 1$ results in $L = K$, and our algebras are exactly the two-power cyclotomics popular in ring cryptography. For roughly cryptographically relevant dimensions, one could consider setting for example $n = 3$ and $m = 7$, resulting in an algebra of dimension 3 over a field L of dimension $6 \cdot 64 = 384$, for a total dimension of 1152. Explicit choices of F can be found in [13]; it is constructively chosen as a subfield of a prime cyclotomic extension of M by considering the Galois group of the extension. It is then straightforward to compute L.

One drawback of the above construction is that we are restricted to odd module dimensions, but if one wanted to consider e.g. $n = 4$ we could replace E with any odd degree extension of M (so that $F \cap E = M$ by coprimality), for example by first setting $E' = M(\zeta_p)$ for prime $p = 3 \mod 4$, then setting $E = M(\zeta_p + \zeta_p^{-1})$, an extension of M of odd degree $\frac{p-1}{2}$. This results in a family of extensions of degree $[E : \mathbb{Q}] = (\frac{p-1}{2}) \cdot \varphi(n)$ and $[EF : E] = n$, corresponding to an even module dimension[2]. Of course, if n is even and not a power of two we can not begin with the construction of [13]; our construction is restricted to prime or prime power n. We consider it likely that this method can be adapted to construct other suitable algebras, but leave finding more algebras for cryptography as an open problem. There are many other methods for finding these objects, for example using local-global theory.

4 Hardness of Search CLWE

For the remainder of this paper, we will always be working in an extension of number fields L/K, where $[L : \mathbb{Q}] = [L : K] \cdot [K : \mathbb{Q}] = n \cdot k$. Recall from the motivation of structured MLWE and the sample algebras given that in practice we seek asymptotic security in k, since the parameter n corresponds to the typically small module dimension. Nonetheless, when considering the comparison to modules, our number fields have dimension nk. We abbreviate the condition $|\sigma_i(\gamma)| = 1$ for all i by $|\gamma| = 1$, since in fact these are equivalent for algebraic γ.

Definition 21. *We define the family of error distributions Σ_α as the set of all Gaussian distributions D_Σ over $\bigoplus_{i=0}^{n-1} u^i L_\mathbb{R}$ with covariance matrix obtained as the distribution of the error in Lemma 12.*

This is the family of error distributions we will claim hardness of search CLWE for; although specifying this family of matrices precisely is not simple, we demonstrate how the error is obtained in the BDD transformation step. For now, we remark that it is a Gaussian distribution whose marginals are Gaussian with variance at most α.

In the following theorem we denote by $\mathcal{A}$−DGS_ξ the problem of sampling a discrete Gaussian $D_{\mathcal{I},\xi}$, where $\mathcal{I}$ is some ideal of the natural order Λ.

[2] Here $\varphi(n)$ denotes the usual Euler totient function.

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

From this we deduce the following corollary, similarly to [14], since the lattice structure of our algebra is merely a special case of their modules.

Corollary 1. *Let $\mathcal{A}, \Lambda, \alpha$ 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 $\sqrt{8Nn} \cdot \xi = (\omega(\sqrt{nk})/\alpha)$.*

The following theorem is our analogy of Lemma 4.10 of [14]. We use $N := k \cdot n^2$ to denote the dimension of our algebra.

Theorem 2. *Given an oracle that solves $CLWE_{q,\Sigma_\alpha}$ for an input $\alpha \in (0,1)$, an integer $q \geq 2$, an order ideal $\mathcal{I} \subset \Lambda$, a number $r \geq \sqrt{2}q \cdot \eta(\mathcal{I})$ 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 $D_{\mathcal{I},r'}$.*

We can then prove Theorem 1 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 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 [26] to sample $D_{\mathcal{I},r}$ on the module representation $\bigoplus_{i=1}^{n-1} u^i L_{\mathbb{R}}$. As usual, we obtain Theorem 2 in two steps, first the main reduction of Lemma 12, then the following quantum step adapted from [26]. We use a form of $\mathcal{A}-\mathrm{BDD}_{L,\delta}$ from [14] where we bound the offset in the norm $\|e\|_{2,\infty} := \max_j \sqrt{(\sum_{i=1}^{n} |\sigma_j(e_i)|^2)} \leq \delta$, where σ denotes the canonical embedding of L.

Lemma 10. *There is an efficient quantum algorithm that given any $N = k \cdot n^2$ dimensional matrix lattice $\mathcal{L} := \sigma_{\mathcal{A}}(\mathcal{I})$ for some ideal $\mathcal{I}$, a real $\delta < \lambda_1(\mathcal{L}^*)/(2\sqrt{2nk})$, and an oracle that solves $\mathcal{A}$-$BDD_{\mathcal{L}^*,\delta}$ with all but negligible probability, outputs an independent sample from $D_{\mathcal{L},\sqrt{n}\omega(\sqrt{\log(nk)})/\sqrt{2}\delta}$.*

4.1 BDD to Search CLWE

We begin with the cyclic algebra analogy of the BDD-to-LWE samples transformation from Section 4 of [16]. As is standard for LWE security, we use the following 'modulo q' definition of BDD:

Definition 22. *For any $q \geq 2$ the $q\mathcal{A}-BDD_{\mathcal{I},d}$ problem is as follows: given an instance of the $\mathcal{A}-BDD_{\mathcal{I},\delta}$ problem $y = x + e$ with solution $x \in \mathcal{I}$ and error $e \in \bigoplus_{i=0}^{n-1} u^i L_{\mathbb{R}}$ satisfying $\|e\|_{2,\infty} \leq \delta$, output $x \mod q\mathcal{I}$.*

We use (a special case of) Lemma 3.5 from [26], which lifts immediately since it is lattice preserving.

Lemma 11. *For any $q \geq 2$ there is a deterministic polynomial time reduction from $\mathcal{A}-BDD_{\mathcal{I},d}$ to $q\mathcal{A}-BDD_{\mathcal{I},d}$.*

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 [16], but has some technical differences induced by the structure of cyclic algebras.

Lemma 12. *Let $\mathcal{A}$ be as in Theorem 1. There is a probabilistic polynomial time algorithm that on input a prime integer $q \geq 2$, a fractional order ideal $\mathcal{I}^\vee \subset \Lambda$, a $q\mathcal{A}-BDD_{L,\alpha q\cdot\omega(\sqrt{\log(nk)})/\sqrt{2nk}\cdot r}$ instance $y = x + e$ where $x \in \mathcal{I}^\vee$ is uniformly random, 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, j^{th} 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. Recall that in the definition of CLWE, a sample is in the form $(a, b) = (a, (a \cdot s)/q + e \mod \Lambda^\vee)$, where e is taken from an error distribution $\psi \in \Sigma_\alpha$.

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

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

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 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 $(z \cdot x)/q = (z \cdot t \cdot x)/q = (a \cdot s)/q \mod \Lambda^\vee$ 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, j^{th} 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

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

where the sum $j + k$ is taken modulo n and the functon $(1 - (1 - \gamma)\mathbb{1}_{j+k \geq n})$ is 1 if $j + k < n$ and γ otherwise[3]. 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 [14]. □

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 n times, and so by symmetry each element of $\sigma_{\mathcal{A}}(e'')$ has non-zero correlation with at most n^2 other entries. Hence, a proportion of at most $\frac{kn^3}{k^2n^4} = \frac{1}{nk}$ of entries of Σ are non-zero.

5 Search To Decision Reduction

In this section 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 [16] and [14] 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 Lemma 4 from [19].

Lemma 13. *Let Λ be the natural order of a cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ and let $\mathcal{I}$ be an ideal of $\mathcal{O}_K$ which splits completely as $\mathcal{I} = \mathfrak{q}_1 ... \mathfrak{q}_k$ as an ideal of $\mathcal{O}_K$. Then, we have the isomorphism*

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

where $\mathcal{R}_i = \bigoplus_{j=0}^{n-1} u^j(\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L)$ 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 $u^n = \gamma + \mathfrak{q}_i$.

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.

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

As in [16], 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. In particular, we restrict our reductions to the family of algebras $\mathcal{A}$ in Section 3.1, but they would be valid in the slightly more general case where the modulus splits completely in the base field K. Now $K = \mathbb{Q}(\zeta_{2^m \cdot n})$ and it follows that if $q \equiv 1 \mod 2^m \cdot n$ then q splits completely into a product of prime ideals $\mathfrak{q}_1, ..., \mathfrak{q}_k$ as an ideal of $\mathcal{O}_K$. Hence, we obtain the decomposition

$$\Lambda / q\Lambda \cong R_1 \times ... \times R_k$$

where R_i is as is Lemma 13.

Also as in [16], we see no way to avoid randomizing the error distribution in the resulting decision problem. However, we face a new issue relating to the automorphisms of $\mathcal{A}$, or lack thereof. To solve this, we require an additional assumption on oracles for the decision CLWE problem. Namely, we assume that an oracle for $\text{CDLWE}_{q,\Upsilon_\alpha}$ 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 assumption implies that for fixed L and K as in Section 3.1 the hardness of the CDLWE problem is invariant under the choice of root of unity γ, and will be required for Lemma 15. We view this as a natural assumption, since the respective natural orders in $\mathcal{A}$ and $\mathcal{A}'$ consist of the same elements and have the same density in their respective algebras. Furthermore, there exist suitable isomorphisms sending $\mathcal{A}$ to $\mathcal{A}'$, which we will define shortly. Consequentially, the main theorem of this section is as follows; we emphasize that our algorithm is only intended to be efficient in the dimension of the base field k, since we expect to fix n as a small constant in practice. Of course, the dimensions of K and L are tied together, so formally to increase $k = 2^{m-1} \cdot (n-1)$ while fixing n we increase the parameter m.

Definition 23. *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 $D_{\boldsymbol{r}}$, where each $r_i := \alpha((n^2 \cdot k)^{1/4} \cdot \sqrt{y_i})$ for $y_1, ..., y_{n^2 \cdot k}$ sampled from $\Gamma(2, 1)$.*

Theorem 3. *Let Λ, q, L, K be as above with $q \in poly(k)$ 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 $CDLWE_{q,\Upsilon_\alpha}$ which runs in time polynomial in k.*

We will prove Theorem 3 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 14). Then, we use a hybrid distribution to define a decision problem in R_i, for which we demonstrate a search to decision reduction (Lemma 15). We then use a hybrid argument to conclude the proof (Lemma 17).

5.1 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 13.

Definition 24. *The $R_i-CLWE_{q,\Sigma_\alpha}$ 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 $u'^n = \alpha_i(\gamma)$ and $xu' = u'\theta(x)$ for $x \in L$. By the construction of K from [13], $\alpha_i(\gamma)$ is a non-norm element since it is some n^{th} 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 n 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 $R_i-\text{CLWE}_{q,\Sigma_\alpha}$ 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 13) 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 14. *There is a deterministic polynomial time reduction from $CLWE_{q,\Sigma}$ to $R_i-CLWE_{q,\Sigma}$.*

Proof. Let $\mathcal{O}_i$ be an oracle for the $R_i-\text{CLWE}_{q,\Sigma}$ problem. Since Lemma 13 defines an isomorphism, it is sufficient to use $\mathcal{O}_i$ to solve the $R_j-\text{CLWE}_{q,\Sigma}$ 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 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 Σ'. 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. □

5.2 Hybrid CLWE and Search-Decision

For this section we must introduce the cyclic algebra analog of the Hybrid LWE distribution used in [16]; we use the decomposition into the rings R_i rather than the Chinese Remainder Theorem.

Definition 25. *For a secret $s \in \Lambda_q^\vee$, distribution Σ over $\bigoplus_j u^j L_\mathbb{R}$, and $i \in [k]$, we define a sample from the distribution $\Pi^i_{q,s,\Sigma}$ over $\Lambda_q \times (\bigoplus_{i=0}^{n-1} u^i L_\mathbb{R})/\Lambda^\vee$ by taking $(a,b) \leftarrow \Pi_{q,s,\Sigma}$ and $h \in \Lambda_q^\vee$ which is uniformly random and independent* $\bmod R_j, j \leq i$ *and* $0 \mod R_j, j > i$*, and outputting $(a, b + h/q)$. If $i = 0$,we define $\Pi^0_{q,s,\Sigma} = \Pi_{q,s,\Sigma}$.*

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

Definition 26. *For $i \in [k]$ and a family of distributions Σ_α, the W-CDLWE$^i_{q,\Sigma_\alpha}$ problem is defined as the problem of finding j given access to $\Pi^j_{q,s,\Sigma}$ for $j \in \{i-1, i\}$ and valid CLWE secret and error distribution s, Σ.*

Lemma 15. *For any $i \in [k]$ there is a probabilistic polynomial-time reduction from R_i-CLWE$_{q,s,\Sigma_\alpha}$ to W-CDLWE$^i_{q,s,\Sigma}$.*

Proof. We follow the standard search-decision methodology of guessing the value of the secret $\bmod R_i$ and then modifying the samples so that the decision oracle tells us whether or not our guess was correct. Note that there are only $N_L(\mathfrak{q}_i\mathcal{O}_L)^n$ possible values of $s \mod R_i$, which is bounded above by q^{n^2}, polynomial in k, 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

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

where $v \in \Lambda_q$ is uniformly random $\bmod R_i$ and $0 \mod R_j$ for $j \neq i$ and $h \in \Lambda_q^\vee$ is uniformly random and independent $\bmod R_j, j < i$ 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 $\bmod R_i$ 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 3. This is the only stage of the proof which enforces that the asymptotic complexity scales only with k and not with n, since we are forced to guess all of $s \bmod R_i$ at once. We do not see a way of separating the coordinates further.

5.3 Worst-Case to Average-Case Decision Reduction

Definition 27. *For $i \in [k]$ and a distribution Υ_α over possible error distributions, an algorithm solves the $CDLWE^i_{q,\Upsilon_\alpha}$ 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 16. *For any $\alpha > 0$ and $i \in [k]$ there is a randomized polynomial-time reduction from W-$CDLWE^i_{q,\Sigma_\alpha}$ to $CDLWEq, {\Upsilon_\alpha}^i$.*

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 [16], 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 4. This choice of Υ_α means that our decision problem is closer to diagonal than 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 conclude with our hybrid argument. We must first show that $\Pi^k_{q,s,\Sigma}$ 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 [21], Lemma 2.4.

Lemma 17. *Let Υ_α be as above and let $s \in \Lambda_q^\vee$. Then given an oracle $\mathcal{O}$ which solves the $CDLWE_{q,\Upsilon_\alpha}$ problem there exists an efficient algorithm that solves $CDLWE^i_{q,\Upsilon_\alpha}$ for some $i \in [k]$ using $\mathcal{O}$.*

Proof. The proof is identical to the ring case, Lemma 5.14 of [16], except that the indexing set $\mathbb{Z}_m^*$ is replaced by k □

6 CLWE Cryptosystem

In this section we present a proof of concept cryptosystem using CLWE, although we do not handle the technical details. To demonstrate our comparison against MLWE our scheme will closely resemble the typical 'compact' LWE cryptography schemes over modules, in particular Kyber (see [6]), although it is likely that an adaptation of Regev style encryption from [26] would suit CLWE as well.

6.1 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. [22]). Finally, we require hardness of a 'normal' form for the CLWE distribution, where s is sampled from the same distribution as the noise e.

Lemma 18. *As long as a non-negligible proportion of elements of Λ_q are invertible there is a polynomial time reduction from the CLWE problem with uniformly random secret s and error distribution Σ to the CLWE problem with secret $s' \leftarrow \Sigma$.*

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 Σ. Sample pairs $(a, b) \leftarrow \Pi_{q,s,\Sigma}$ 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 by assumption a non-negligible fraction of elements of Λ_q are invertible, with high probability this step takes polynomial time.

Now, given a pair $(a_i, b_i) \leftarrow \Pi_{q,s,\Sigma}$, we obtain a sample from the CLWE distribution $\Pi_{q,e_1,\Sigma}$ 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 Σ. □

Recall that for the decision problem, we are interested in asymptotic complexity in k. For our choice of number fields from Section 3.1, Propositions 1 and 4 of [19] give us that Λ_q is isomorphic to a direct product of k matrix algebras of dimension n over $\mathbb{Z}_q$, for which a non-negligible proportion of elements are invertible. Combining these properties, the hardness of the decision CLWE problem over $\Lambda_q \times \Lambda_q$,where a is uniformly random and $s, e \leftarrow \Sigma$ for some discretized error distribution Σ, follows.

6.2 Sample Cryptosystem

Our scheme is parameterized by an algebra $\mathcal{A} := (L/K, \theta, \gamma)$, where $\mathcal{A}$ is as in Section 3.1, an error distribution Σ, and a prime modulus $q \equiv 1 \mod (2^m \cdot n)$ (recall $K = \mathbb{Q}(\zeta_{2^m \cdot n})$). We will denote with bold faced letters the vector form of an element of Λ_q, e.g. if $a = a_0 + u a_1 + ... + u^{n-1} a_{n-1}$ then $\mathbf{a} = (a_0, a_1, ..., a_{n-1})$. We note that $\mathcal{O}_L / q\mathcal{O}_L$ has a polynomial representation of dimension $n \cdot k$, and so we encode our message $m \in \{0,1\}^{k \cdot n^2}$ as an entry of Λ_q as a vector $\mathbf{m}$ of n $\{0,1\}$ 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 $e \leftarrow \Sigma$, and outputs public key $\mathbf{a}, \mathbf{b}$.
– To encrypt $\mathbf{m} \in \{0,1\}^{k \cdot n^2}$, 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 5. There are two benefits of instantiating this scheme in the cyclic algebra setting rather than over modules as in [6], both following from the matrix embedding ϕ. Firstly, in the module setting Alice must publish a matrix $\mathbf{A}$ rather than the vector $\mathbf{a}$ in her key, since $\phi(a)$ lets us generate a matrix; this saves a factor of n 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 n.

Example 2. recall our sample parameters from Section 3.1 with $[K : \mathbb{Q}] = 128$ and $[L : K] = 3$, such that $[\Lambda : \mathbb{Z}] = 1152$. 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 3 over a ring L of dimension 384. Our public key $(\mathbf{a}, \mathbf{b})$ requires storing only 6 elements of $R_q = O_L/q \cdot O_L$ rather than 12 in the form $(A, \mathbf{b})$ and our message consists of 1152 bits, corresponding to the total dimension of the algebra rather than the module versions 384 which corresponds to the field dimension. Our ciphertext sizes are the same. 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 1152. Both public keys are in the form (a, s) and so require equivalent levels of storage (6 elements of a field of dimension 384 or 2 in dimension 1152), and the same phenomenon is true of ciphertext sizes and message length. Hence, we hope to gain in security without losing efficiency.

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 $\mathcal{A} = (L/K, \theta, \gamma)$ 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 are 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 isomorphic 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 $(x_0, \gamma\theta(x_{n-1}), \gamma\theta^2(x_{n-2}), \ldots, \gamma\theta^{n-1}(x_1))$, by setting $a' = x_0 + u\gamma\theta(x_{n-1}) + \cdots + u^{n-1}\gamma\theta^{n-1}(x_1)$ and observing that θ^n 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 [1] or Kyber [6]. 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 the challenge ciphertext $\mathbf{v}$ with $\mathbf{v}'$ by replacing b with b'. 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 $\mathbf{v}'$ 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 $k \cdot n^2$.*

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

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

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.

Acknowledgements

Grover was supported by an ACE-CSR PhD grant. The authors would like to thank Prof. Damien Stehlé and Dr. Roope Vekhalahti for many fruitful discussions.

References

1. Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange—a new hope. In: 25th USENIX Security Symposium (USENIX Security 16). pp. 327–343 (2016)
2. Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. Mathematische Annalen **296**(1), 625–635 (1993)
3. Baumslag, G., Fazio, N., Nicolosi, A.R., Shpilrain, V., Skeith III, W.E.: Generalized learning problems and applications to non-commutative cryptography. In: Provable Security, pp. 324–339. Springer (2011)
4. Biasse, J.F., Song, F.: 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)
5. Bootland, C., Castryck, W., Vercauteren, F.: On the Security of the Multivariate Ring Learning with Errors Problem (2018), published: Cryptology ePrint Archive, Report 2018/966
6. Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehlé, D.: CRYSTALS-Kyber: a CCA-secure module-lattice-based KEM. In: 2018 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 353–367. IEEE (2018)
7. Campbell, P., Groves, M., Shepherd, D.: Soliloquy: A cautionary tale (2015)
8. Caruso, X., Le Borgne, J.: Fast multiplication for skew polynomials. In: Proceedings of the 2017 ACM on International Symposium on Symbolic and Algebraic Computation. pp. 77–84. ACM (2017)
9. Cheng, Q., Zhuang, J.: LWE from Non-commutative Group Rings. arXiv preprint arXiv:1612.06670 (2016)
10. Cramer, R., Ducas, L., Peikert, C., Regev, O.: 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)
11. Cramer, R., Ducas, L., Wesolowski, B.: 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)
12. Jozsa, R.: Quantum factoring, discrete logarithms, and the hidden subgroup problem. Computing in Science & Engineering **3**(2), 34–43 (2001)
13. Lahtonen, J., Markin, N., McGuire, G.: Construction of Multiblock Space–Time Codes From Division Algebras With Roots of Unity as Nonnorm Elements. IEEE Transactions on Information Theory **54**(11), 5231–5235 (Nov 2008)
14. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. Designs, Codes and Cryptography **75**(3), 565–599 (2015)
15. Luzzi, L., Vehkalahti, R., Ling, C.: Almost universal codes for MIMO wiretap channels. IEEE Transactions on Information Theory **64**(11), 7218–7241 (2018)
16. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 1–23. Springer (2010)

17. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. SIAM Journal on Computing **37**(1), 267–302 (2007)
18. Oggier, F., Belfiore, J.C., Viterbo, E.: Cyclic division algebras: A tool for space-time coding. Now Publishers Inc (2007)
19. Oggier, F., A. Sethuraman, B.: Quotients of Orders in Cyclic Algebras and Space-Time Codes. Advances in Mathematics of Communications **7** (2012)
20. Pedrouzo-Ulloa, A., Troncoso-Pastoriza, J.R., Pérez-González, F.: On Ring Learning with Errors over the Tensor Product of Number Fields. arXiv preprint arXiv:1607.05244 (2016)
21. Peikert, C.: An efficient and parallel Gaussian sampler for lattices. In: Annual Cryptology Conference. pp. 80–97. Springer (2010)
22. Peikert, C.: How (not) to instantiate ring-LWE. In: International Conference on Security and Cryptography for Networks. pp. 411–430. Springer (2016)
23. Peikert, C., Regev, O., 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. pp. 461–473. ACM (2017)
24. Pierce, R.S.: Cyclic Division Algebras. In: Associative Algebras, pp. 276–293. Springer New York, New York, NY (1982)
25. Puchinger, S., Wachter-Zeh, A.: Fast operations on linearized polynomials and their applications in coding theory. Journal of Symbolic Computation **89**, 194–215 (2018)
26. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. Journal of the ACM (JACM) **56**(6), 34 (2009)
27. Reiner, I.: Maximal orders. L.M.S. monographs, Academic Press (1975)
28. Vehkalahti, R., Hollanti, C., Lahtonen, J., Ranto, K.: On the densest MIMO lattices from cyclic division algebras. IEEE Transactions on Information Theory **55**(8), 3751–3780 (2009)

A Regarding BCV Style Attacks

For $m = 2$, the m-RLWE protocol of [20] 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 modulo R_q. The example use case given in [20] 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 [5] 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, [5] 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 $[E : K] = n$, 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 Section 15.1, Proposition a of [24]). 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 [5] does not threaten our algebraic structure.