Middle-Products of Skew Polynomials and Learning with Errors

Cong Ling and Andrew Mendelsohn

Department of EEE, Imperial College London, London, SW7 2AZ, United Kingdom.
andrew.mendelsohn18@imperial.ac.uk
c.ling@imperial.ac.uk

Abstract. We extend the middle product to skew polynomials, which we use to define a skew middle-product Learning with Errors (LWE) variant. We also define a skew polynomial LWE problem, which we connect to Cyclic LWE (CLWE), a variant of LWE in cyclic division algebras. We then reduce a family of skew polynomial LWE problems to skew middle-product LWE, for a family which includes the structures found in CLWE. Finally, we give an encryption scheme and demonstrate its IND-CPA security, assuming the hardness of skew middle-product LWE.

Keywords: middle product · LWE · cyclic division algebras · skew polynomials

1 Introduction

The development of efficient quantum algorithms for cryptographic problems (e.g. [21]) has lead to the development of *post*-quantum cryptography, which relies on computationally intractable problems for both classical and quantum computers. A prime candidate for a family of such computationally intractable problems are *lattice* problems, following the pioneering work of Ajtai [1]. In particular, much post-quantum cryptographic functionality is based on the Learning with Errors (LWE) problem, introduced by Regev [18].

LWE-style problems consist of solving systems of noisy linear equations. Over the integers, LWE loosely asks a challenger to find $\mathbf{s} \in \mathbb{Z}_q^n$ from a number of samples of the form $(\mathbf{a}_i, \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i \bmod q)$, where $\mathbf{a}_i \in \mathbb{Z}_q^n$ and e_i is some noise. However, cryptosystems based on LWE have sub-optimal storage requirements and computation with LWE samples is often inefficient, due to the relative inefficiency of high-dimensional matrix multiplication. For this reason, structured variants of LWE have been introduced.

These include Ring LWE (RLWE) [15], which uses multiplication in the ring of integers of a number field to create multiple correlated LWE samples. For instance, if $\mathcal{R}$ is the ring of integers of the $2n$th cyclotomic field for power-of-two n, then $\mathcal{R} = \mathbb{Z}[x]/(x^n + 1)$ and multiplication on a fixed basis by a polynomial

$a \in \mathcal{R}$ can be represented by a matrix

$$\begin{pmatrix} a_0 & -a_{n-1} & \cdots & -a_1 \\ a_1 & a_0 & \dots & -a_2 \\ \vdots & \vdots & \ddots & \vdots \\ a_{n-1} & a_{n-2} & \dots & a_0 \end{pmatrix}.$$

Other structured forms of LWE have been studied, such as PLWE [22], which considers $\mathcal{R} = \mathbb{Z}[x]/(f(x))$ for a broader range of $f(x)$, and CLWE [9], which developed LWE from orders in cyclic division algebras (CDAs). These variants both use algebraic objects which permit matrix representations over $\mathbb{Z}$ to rewrite multiplication by an element a as multiplication by an integral matrix.

Another variant, middle-product LWE (MPLWE) [19], replaced ring multiplication with the *middle product*, denoted $\odot$. This product takes two polynomials a, b and outputs a polynomial whose coefficients are the 'middle' coefficients of the product $a \cdot b$, discarding higher and lower order terms. In particular, given $a = \sum_{i=0}^{d_a-1} x^i a_i$, $b = \sum_{i=0}^{d_b-1} x^i b_i$ with $d_a + d_b - 1 = d + 2k$ for some d, k, we have

$$a \odot_d b = \left\lfloor \frac{(a \cdot b) \bmod x^{k+d}}{x^k} \right\rfloor.$$

The discarding of coefficients allows for fast algorithms to compute middle products [10], [8] and this product has a matrix presentation such that samples of shape $(a, a \odot r + e)$ form structured instances of LWE. In particular, one can write

$$a \odot_d r = \begin{pmatrix} a_0 & a_1 & a_2 & \dots & a_{d_a-1} & 0 & \dots & 0 \\ 0 & a_0 & a_1 & \dots & a_{d_a-2} & a_{d_a-1} & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \dots & \vdots \\ 0 & \dots & \dots & \dots & 0 & a_0 & \dots & a_{d_a-1} \end{pmatrix} \cdot \begin{pmatrix} r_{d_r-1} \\ \vdots \\ r_1 \\ r_0 \end{pmatrix}$$

In [19] a reduction from a family of PLWE problems to MPLWE was given, guaranteeing that MPLWE is at least as hard as the hardest PLWE problem in the family. Notably the chosen family includes RLWE instances. They also gave a public key encryption scheme and proved its IND-CPA security, assuming hardness of MPLWE.

Our Contribution We develop a novel form of MPLWE for skew polynomial rings, which are a noncommutative form of polynomial ring, named 'Skew MPLWE' (SMPLWE). We define the middle product for such rings and also a novel structured form of LWE for skew polynomial rings, named 'skew polynomial LWE' (SPLWE). We show that this LWE variant includes CLWE instances, reduce a family of SPLWE problems to SMPLWE, and give a PKE scheme.

We state four motivations for this work:

1. We define and make use of (to our knowledge) the first structured LWE-variant from skew polynomials. This was implicit in [9], but the connection

was never utilised other than for multiplication algorithms. This appears a promising avenue of future research, given the well-studied properties of skew polynomial rings and their profitable application by coding theorists.

2. We continue the study of LWE in CDAs. Defining SMPLWE and SPLWE and relating them to CLWE provides further indications of the precise security level of CLWE, which is believed to lie somewhere between that of RLWE and MLWE, but more precise understanding is lacking. Our reduction provides new quantitative information on CLWE.
3. SMPLWE enjoys a reduction from a family of SPLWE problems (including CLWE-style problems). This provides SMPLWE with a strong security guarantee and may be preferable in some contexts to CLWE, for this reason.
4. SMPLWE, like MPLWE, enjoys fast multiplication algorithms. Fast algorithms for skew polynomials exist [7], and it seems likely that these could be used to efficiently compute the skew middle product. This yields a cryptographic scheme which is both efficient and, as explained above, secure.

Our reduction holds for a restricted parameter set relative to [19], since it appears the noncommutative structure of our rings means that for only some parameters is SMPLWE structured LWE (in the notation of [19], when $n = m = d$). In more detail, we consider quotients of skew polynomial rings of the form $\mathcal{O}_L[u, \theta]/(u^d - \gamma)$, where L is a number field with ring of integers $\mathcal{O}_L$, K is an index d subfield of L such that $\mathrm{Gal}(L/K)$ is generated by an automorphism θ, u satisfies $ux = \theta(x)u$ for any $x \in \mathcal{O}_L$, and $\gamma \in \mathcal{O}_K$, and prove our results for middle product samples $(a, a \odot_d r + e)$, where $\deg(a) = d - 1$ and $\deg(r) = 2(d - 1)$. In this setting, we set $a \odot_d r = \left\lfloor \frac{(a \cdot r) \bmod u^{2d-1}}{u^{d-1}} \right\rfloor$ and can write

$$a \odot_d r = \begin{pmatrix} a_{d-1} & \theta(a_{d-2}) \dots & \theta^{d-1}(a_0) & 0 & \dots & 0 \\ 0 & \theta(a_{d-1}) \dots & \theta^{d-1}(a_1) & \theta^d(a_0) & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots \theta^{d-1}(a_{d-1}) & \dots & \theta^{d-3}(a_1) & \theta^{d-2}(a_0) \end{pmatrix} \begin{pmatrix} r_0 \\ r_1 \\ \vdots \\ r_{d_r-1} \end{pmatrix}$$

We then define two problems: $\mathrm{SPLWE}_{q,s,f,\chi}$ is the problem of distinguishing samples of the form $(a_i, a_i s + e_i \bmod q)$ from samples uniform over the domain, and $\mathrm{SMPLWE}_{q,s,d,d,\chi'}$ is the challenge of distinguishing samples of the form $(a_i, a_i \odot_d s + e_i)$ from those uniform over the domain, where a_i and s are skew polynomials of bounded degree and e_i is added noise. We then prove

Main Reduction (Theorem 1)**.** Let $d > 0, q \geq 2$, and χ an error distribution. Then there exists a ppt. reduction from $\mathrm{SPLWE}_{q,s,f,\chi}$ for any polynomial $f(u) = u^d - \gamma \in \mathcal{O}_L[u, \theta]$ with $\gamma \in \mathcal{O}_K \setminus \{0\}$ coprime with q, to $\mathrm{SMPLWE}_{q,s,d,d,\chi'}$.

This result reduces a family of SPLWE problems to SMPLWE - a family which includes CLWE-style instances. To achieve this, new families of linear transformations on coefficients of skew polynomials are introduced. We then give a PKE scheme and demonstrate its IND-CPA security, if SMPLWE is hard.

We note here that we consider a family of SPLWE problems under the coefficient embedding. These SPLWE problems include the ones considered in CLWE, but in that setting the canonical embedding was used. It is not currently clear what the relationship between CLWE in the coefficient and in the canonical embedding is, but it seems likely that, in a similar way as holds for RLWE, CLWE under the coefficient embedding is still a 'hard' problem, although we stress that we do not have any formal proofs of the security of CLWE under the coefficient embedding. However, we note the work of [20] and consider it reasonable to suggest that CLWE in the canonical and coefficient embeddings can be related via a linear transformation with limited loss in parameter quality. We provide evidence toward this end in Appendix B.

Prior Work and Paper Organisation MPLWE was introduced in [19] and CLWE in [9]. More on middle product-based cryptography can be found in [14], [4], [5], [23]. In [17], MPLWE was related to a number of LWE variants, such as RLWE. We note the extensive use of skew polynomials in coding theory [3].

Preliminaries are in Section 2, we recollect LWE in Section 3, skew polynomials in Section 4, and CDAs in Section 5. We introduce the skew middle product in Section 6, give a reduction from SPLWE to SMPLWE in Section 7, provide a PKE scheme in Section 8, and then conclude.

2 Preliminaries

If $\mathbf{v}$ is an n-dimensional vector, we denote by $\bar{\mathbf{v}}$ the n-dimensional vector whose entries are those of $\mathbf{v}$ in reverse order; i.e. if $\mathbf{v} = (v_1, ..., v_n)^T$, then $\bar{\mathbf{v}} = (v_n, ..., v_1)^T$. We prove IND-CPA security of our cryptosystem below. Recall:

Definition 1. ([12]) Let $\Pi = (\text{Gen, Enc, Dec})$ be a PKE scheme, and $\mathcal{A}$ be an adversary. We say Π is *indistinguishable under chosen-plaintext attack* if any ppt. adversary in the following experiment $\text{PubK}_{\mathcal{A},\Pi}(n)$ has negligible advantage:

1. Gen is run to obtain keys (pk, sk).
2. Adversary $\mathcal{A}$ is given pk, and outputs a pair of equal-length messages m_0, m_1 in the message space.
3. A uniform bit $b \in \{0, 1\}$ is chosen, and then a ciphertext $c \leftarrow \text{Enc}_{pk}(m_b)$ is computed and given to $\mathcal{A}$. We call c the challenge ciphertext.
4. $\mathcal{A}$ outputs a bit b'. The output of the experiment is 1 if $b' = b$, and 0 otherwise. If $b' = b$ we say that $\mathcal{A}$ succeeds.

That is, $\Pr[\text{PubK}_{\mathcal{A},\Pi}(n) = 1] \leq \frac{1}{2} + \text{neg}(n)$.

To complete the proof, we will rely on properties of hash functions:

Definition 2. A family $\mathcal{H}$ of hash functions $h : X \to Y$ of finite cardinality is called *universal* if $\Pr_{h \leftarrow U(\mathcal{H})}[h(x_1) = h(x_2)] = 1/|Y|$, $\forall\, x_1 \neq x_2 \in X$.

The *statistical distance* between two distributions D, D' over a discrete set S is defined $\Delta(D, D') = \frac{1}{2}\sum_{x \in S} |D(x) - D'(x)|$. The uniform distribution over a finite set S' is denoted $U(S')$.

Lemma 1. *[19, Lemma 2.1] Let X, Y, Z be finite sets. Let $\mathcal{H}$ be a universal hash function family $h : X \to Y$ and $f : X \to Z$ be arbitrary. Then for any random variable T taking values in X, and $\gamma(T) = \max_{t \in X} \Pr[T = t]$, we have:*

$$\Delta((h, h(T), f(T)), (h, U(Y), f(T))) \leq \frac{1}{2} \cdot \sqrt{\gamma(T) \cdot |Y| \cdot |Z|}$$

3 Learning with Errors and Middle Products

The Middle Product The middle product can be thought of as the multiplication rule which takes two polynomials, multiplies them together, then discards the lower and higher coefficients, forming a polynomial whose coefficients are the 'middle' part of the product. Formally, if $\mathcal{R}$ is an arbitrary ring and $\mathcal{R}^{<d}[x]$ denote the polynomials over $\mathcal{R}$ of degree at most $d-1$:

Definition 3. Let $d_a, d_b, d, k \in \mathbb{N}$ such that $d_a + d_b - 1 = d + 2k$. The middle-product of $a \in \mathcal{R}^{<d_a}[x]$ and $b \in \mathcal{R}^{<d_b}[x]$ is defined

$$\begin{aligned} \odot_d : & \mathcal{R}^{<d_a}[x] \times \mathcal{R}^{<d_b}[x] \to \mathcal{R}^{<d}[x], \\ & (a, b) \mapsto a \odot_d b = \left\lfloor \frac{(a \cdot b) \bmod x^{k+d}}{x^k} \right\rfloor . \end{aligned}$$

We can now define middle product learning with errors, following [19]:

Definition 4. (MPLWE distribution) Let $n, d > 0, q \geq 2$, and χ be a distribution over $\mathbb{R}^{<d}[x]$. For $s \in \mathbb{Z}_q^{<n+d-1}[x]$, define the distribution $\mathrm{MP}_{q,n,d,\chi}(s)$ over $\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x]$ as the distribution obtained by sampling $a \leftarrow U\left(\mathbb{Z}_q^{<n}[x]\right), e \leftarrow \chi$ and outputting $(a, b = a \odot_d s + e)$.

Definition 5. (decision MPLWE) Let $n, d > 0, q \geq 2$, and χ be a distribution over $\mathbb{R}^{<d}[x]$. Then the decision MPLWE problem, $\mathrm{MPLWE}_{n,d,q,\chi}$, consists in distinguishing between arbitrarily many samples from $\mathrm{MP}_{q,n,d,\chi}(s)$ and the same number of samples from $U\left(\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x]\right)$, with non-negligible probability over $s \leftarrow U\left(\mathbb{Z}_q^{<n+d-1}[x]\right)$.

4 Skew Polynomials

A skew polynomial ring over a field is defined as follows:

Definition 6. Let $\mathbb{F}$ be a field and θ be an automorphism of $\mathbb{F}$. Then $\mathbb{F}[u, \theta] := \{\sum_{i=0}^{n} u^i x_i : x_i \in \mathbb{F}\}$, the set of polynomials in u with coefficients in $\mathbb{F}$ equipped with standard polynomial addition and having polynomial multiplication subject to the condition $xu = u\theta(x)$ for all $x \in \mathbb{F}$, is called a skew polynomial ring.

The multiplication rule means that for non-trivial choice of θ, $\mathbb{F}[u, \theta]$ is a non-commutative ring. If $\mathbb{F}^\theta$ is the fixed field of θ, $\mathbb{F}^\theta = \{x \in \mathbb{F} : \theta(x) = x\}$, and θ has order d, then $\mathbb{F}^\theta[u^d]$ is the largest commutative subring of $\mathbb{F}[u, \theta]$. The elements of this subring are called *central* and generate two-sided ideals of $\mathbb{F}[u,\theta]$. For more on skew polynomials, see [16], [11, Chapter 8] or Appendix A.

One may restrict the coefficients to be taken from some subring of a field, and for MPLWE in skew polynomial rings we will indeed restrict the coefficients to the ring of integers of a number field. An important construction of skew polynomial rings (other examples can be found in Appendix A) is the following:

Example 1. Let $L/\mathbb{Q}$ be a finite Galois extension, and $\theta \in \mathrm{Gal}(L/\mathbb{Q})$ with fixed field K, such that $[L:K]=d$ and $\mathrm{Gal}(L/K)$ is cyclic. Then $\mathcal{O}_L[u,\theta]$ is a skew polynomial ring with center $\mathcal{O}_K[u^d]$.

Skew Polynomial Learning with Errors In this section we define a Learning with Errors distribution sampling skew polynomials, and state search and decision problems for that distribution. Below, $R_q := R/qR$ and $R_{q,f} := R/(q,f)R$ for a ring R.

Definition 7. Let $q \geq 2$ and $d \geq 1$. Let θ be an automorphism of L of degree d, $R := \mathcal{O}_L[u,\theta]$, $L_\mathbb{R} := L \otimes \mathbb{R}$, $f \in R$ be a monic central skew polynomial of degree n, and $s \in R_{q,f}$. To obtain a sample from the Skew Polynomial Learning with Errors distribution (SPLWE) $\mathrm{SP}_{q,s,f,\chi}$, sample $a \leftarrow U(R_{q,f})$, $e \xleftarrow{\chi} L_\mathbb{R}[u,\theta]/fR$, and output $(a, as+e \bmod q) \in R_{q,f} \times L_\mathbb{R}[u,\theta]/(q,f)R$.

The decision problem is then defined as follows:

Definition 8. (decision SPLWE) Let Υ be a distribution on a family of error distributions over $L_\mathbb{R}[u,\theta]$, and $U(\cdot)$ be the uniform distribution. The decision SPLWE problem $\mathrm{SPLWE}_{q,s,f,\chi}$ is on input a number of independent samples from either $\mathrm{SP}_{q,s,f,\chi}$ for random $(s,\chi) \leftarrow U(R_{q,f}) \times \Upsilon$ or $U(R_{q,f} \times L_\mathbb{R}[u,\theta]/(q,f)R)$, to decide which is the case with non-negligible advantage.

Useful Matrices for Manipulating Skew Polynomials In this section we will define and prove basic properties of a number of linear transformations on the coefficients of skew polynomials, which we later use in establishing the hardness of SMPLWE and a cryptosystem based off it. We define these as matrices, and specialise to the skew polynomial rings of Example 1. We begin with:

Definition 9. Let $f \in \mathcal{O}_L[u,\theta]$ be a monic central skew polynomial of degree m. Let $a \in \mathcal{O}_L[u,\theta]$. Define $\mathrm{Rot}_f^d(a)$ as the $d \times m$ matrix with ith row given by the coefficients of $a \cdot u^{i-1} \bmod f$, for $i = 1, ..., d$.

It is immediate that if $a \equiv a' \bmod f$, then $\mathrm{Rot}_f^d(a) = \mathrm{Rot}_f^d(a')$. Moreover, $\mathrm{Rot}_f^d(ab) = \mathrm{Rot}_f^d(b)\,\mathrm{Rot}_f^d(a)$. When $m = d$, we will write $\mathrm{Rot}_f(a)$ for $\mathrm{Rot}_f^d(a)$.

Definition 10. Let $f \in \mathcal{O}_L[u,\theta]$ be a monic central skew polynomial of degree m. Define $M_{f,\theta}$ as the $m \times m$ matrix with entries such that $M_{f,\theta} \cdot \mathbf{a}$ has ith entry

$$\left(\left(\sum_{j=1}^{m} u^{i+j-2}\theta^{i-1}(a_{j-1})\right) \bmod f\right) \bmod u.$$

We introduce this matrix for the following reason:

$$\begin{aligned}\left(\sum_{j=1}^{m} u^{i+j-2}\theta^{i-1}(a_{j-1}) \bmod f\right) \bmod u &= \left(\sum_{j=1}^{m} u^{j-1}a_{j-1}u^{i-1} \bmod f\right) \bmod u\\ &= \left(\sum_{j=1}^{m} u^{j-1}a_{j-1}u^{i-1} \bmod f\right) \bmod u\\ &= \left(au^{i-1} \bmod f\right) \bmod u,\end{aligned}$$

which is the constant coefficient of $au^{i-1} \bmod f$, and hence

$$M_{f,\theta} \cdot \mathbf{a} = \mathrm{Rot}_f(a) \cdot (1, 0, ..., 0)^T.$$

Example: Suppose $f(u) = u^d - \gamma$ for some $\gamma \in \mathcal{O}_K$ and $\deg(a) = d-1$. Then

$$\mathrm{Rot}_f(a) = \begin{pmatrix} a_0 & a_1 & ... & a_{d-1} \\ \gamma\theta(a_{d-1}) & \theta(a_0) & ... & \theta(a_{d-2}) \\ \vdots & \vdots & \ddots & \vdots \\ \gamma\theta^{d-1}(a_1) & \gamma\theta^{d-1}(a_2) & ... & \theta^{d-1}(a_0) \end{pmatrix},\ M_{f,\theta} = \begin{pmatrix} 1 & 0 & ... & 0 & 0 \\ 0 & 0 & ... & 0 & \gamma\theta \\ 0 & 0 & ... & \gamma\theta^2 & 0 \\ \vdots & \vdots & ... & \vdots & \vdots \\ 0 & \gamma\theta^{d-1} & ... & 0 & 0 \end{pmatrix}$$

We introduce a kind of generalised Toeplitz matrix which we will later require:

Definition 11. Let $d, k > 0$. Let $r \in \mathcal{O}_L^{<k+1}[u, \theta]$. Set $\mathrm{GToep}^{d,k+1}(r)$ to be the $d \times (k+d)$ matrix whose i, jth entry is given by $\theta^{j-1}(r_{k-j+i})$.

This definition is important for writing the middle product in matrix form, as we shall see later. It also has the following property: if $f(u) = u^d - \gamma$ for some $\gamma \in K$ and $a \in \mathcal{O}_L[u, \theta]^{<d}$ is a skew polynomial, there exists a $2d-1 \times d$ matrix N_f and skew polynomial $\tilde{a}$ such that $\mathrm{GToep}^{d,d}(a) \cdot N_f = \mathrm{Rot}_f(\tilde{a})$. Formally:

Proposition 1. *Let $a \in \mathcal{O}_L[u, \theta]^{<d}$, $f(u) = u^d - \gamma$, and θ have order d. Then there exists a $2d-1 \times d$ matrix N_f and a skew polynomial $\tilde{a}$ such that* $\mathrm{GToep}^{d,d}(a) \cdot N_f = \mathrm{Rot}_f(\tilde{a})$. *Moreover, if $a = a_0 + ua_1 + ... + u^{d-1}a_{d-1}$, we have $\tilde{a} = a_{d-1} + u\theta(a_{d-2}) + ... + u^{d-1}\theta^{d-1}(a_0)$.*

Proof. Write $a = a_0 + ua_1 + ... + u^{d-1}a_{d-1} \in \mathcal{O}_L[u, \theta]$. $\mathrm{GToep}^{d,d}(a)$ has the form

$$\mathrm{GToep}^{d,d}(a) = \begin{pmatrix} a_{d-1} & \theta(a_{d-2}) & \dots & \theta^{d-1}(a_0) & 0 & \dots & 0 \\ 0 & \theta(a_{d-1}) & \dots & \theta^{d-1}(a_1) & \theta^d(a_0) & \dots & 0 \\ & & \ddots & \ddots & \ddots & & \\ 0 & 0 & \dots & \theta^{d-1}(a_{d-1}) & \dots & \theta^{d-3}(a_1) & \theta^{d-2}(a_0) \end{pmatrix}$$

Note that the entries of each column of $\mathrm{GToep}^{d,d}(\cdot)$ all feature the same power of θ. Since $\mathrm{GToep}^{d,d}$ has size $d \times 2d-1$ and Rot_f size $d \times d$, any matrix N such that $\mathrm{GToep}^{d,d} \cdot N = \mathrm{Rot}_f$ must have size $2d-1 \times d$. Setting N_f to be the matrix

$$N_f = \left(\begin{array}{ccccc} \multicolumn{5}{c}{I_d} \\ \hline \gamma & 0 & 0 & ... & 0 \\ 0 & \gamma & 0 & ... & 0 \\ ... & ... & \ddots & ... & ... \\ 0 & 0 & ... & \gamma & 0 \end{array}\right),$$

where I_d is the $d \times d$ identity matrix, one finds

$$\mathrm{GToep}^{d,d}(a) \cdot N_f = \begin{pmatrix} a_{d-1} & \theta(a_{d-2}) & ... & \theta^{d-1}(a_0) \\ \gamma\theta^{d}(a_0) & \theta(a_{d-1}) & ... & \theta^{d-1}(a_1) \\ \vdots & \vdots & ... & \vdots \\ \gamma\theta^{d}(a_{d-2}) & \gamma\theta(a_{d-3}) & ... & \theta^{d-1}(a_{d-1}) \end{pmatrix},$$

which is $\mathrm{Rot}_f(\tilde{a})$, where $\tilde{a} = a_{d-1} + u\theta(a_{d-2}) + ... + u^{d-1}\theta^{d-1}(a_0)$. □

5 Cyclic Division Algebras and CLWE

In this section we review Cyclic LWE. Suppose L/K is a finite Galois extension of number fields of degree d and $\langle\theta\rangle = \mathrm{Gal}(L/K)$. Consider

$$\mathcal{A} := L + uL + ... + u^{d-1}L,$$

where u is such that $\big)$ $u^d = \gamma$ for some $\gamma \in K$ and 2$\big)$ $ux = \theta(x)u$ for all $x \in L$. Then we call $\mathcal{A}$ a cyclic algebra over K, and write $(L/K, \theta, \gamma)$. When $\gamma \in \mathcal{O}_K$, $\mathcal{A}$ contains a discrete subring

$$\Lambda := \mathcal{O}_L + u\mathcal{O}_L + ... + u^{d-1}\mathcal{O}_L.$$

An important property of cyclic algebras is the *division* property; we say a cyclic algebra $\mathcal{A}$ is division if every element has a multiplicative inverse. Division algebras are noncommutative equivalents of fields (and sometimes known as skew fields). The following provides a useful criterion for a cyclic algebra to be division:

Definition 12. An element α of K is *non-norm* if there does not exist an element $x \in L$ such that $\alpha^i = N_{L/K}(x)$, for $0 < i < [L : K]$.

Proposition 2. *[2] The cyclic algebra $\mathcal{A}$ is a division algebra if and only if γ is a non-norm element.*

We connect CDAs with skew polynomial rings via the following:

Lemma 2. *Let $[L : K] = d$ and $\langle\theta\rangle = \mathrm{Gal}(L/K)$. Then $\Lambda \cong \mathcal{O}_L[u, \theta]/(u^d - \gamma)$.*

Proof. We define a map $\varphi : \mathcal{O}_L[u,\theta] \to \Lambda$ via

$$g(u) \mapsto g'(u) := g(u) \bmod (u^d - \gamma) \mapsto g',$$

where $g(u) = g_0 + ug_1 + ... + u^{k-1}g_{k-1}$ is a skew polynomial in $\mathcal{O}_L[u,\theta]$ and $g' \in \Lambda$ has coefficients g'_i, $i = 0, ..., d-1$. This map is surjective, since any element of Λ can be written $g = g_0 + ug_1 + ... + u^{d-1}g_{d-1}$ with coefficients in $\mathcal{O}_L$, so $\varphi(g_0 + ug_1 + ... + u^{d-1}g_{d-1}) = g$ trivially. Let $x \in \ker(\varphi)$, so $\varphi(x) = 0$. This means $g'(u) = 0$, since the second map sends the u^i-coefficients of the skew polynomial to the u^i-coefficients of the element of Λ, so an element of the kernel is in the ideal generated by $u^d - \gamma$ in $\mathcal{O}_L[u,\theta]$. This ideal is two-sided, as $u^d - \gamma$ is central, so $\mathcal{O}_L[u,\theta]/(u^d - \gamma)$ is a ring, and we obtain an isomorphism of rings. □

When $K = \mathbb{Q}(\zeta_m)$ is the mth cyclotomic field, L/K is such that $\mathrm{Gal}(L/K)$ is cyclic, and $\gamma \in \mathcal{O}_K^\times$ with $\gamma \notin N_{L/K}(L^\times)$, then Λ is a maximal order in a CDA [9]. This enables us to connect SPLWE, CLWE and SMPLWE (defined below).

CLWE In [9], an LWE problem was defined in Λ via the CLWE distribution. We state a version in which a and s are sampled from Λ. Below $L_\mathbb{R} := L \otimes \mathbb{R}$.

Definition 13. Let L/K be a Galois extension of number fields with $[L : K] = d$ and $\mathrm{Gal}(L/K)$ cyclic, generated by θ. Let $\mathcal{A} := (L/K, \theta, \gamma)$ be the resulting cyclic K-algebra with element u such that $u^d = \gamma \in \mathcal{O}_K$ and γ satisfying the non-norm condition. Let Λ be the natural order of $\mathcal{A}$. For an error distribution ψ over $\bigoplus_{i=0}^{d-1} u^i L_\mathbb{R}$, $q \geq 2$, and secret $s \in \Lambda_q$, 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 $(a, b) = (a, a \cdot s + e \bmod q\Lambda) \in \left(\Lambda_q, \bigoplus_{i=0}^{d-1} u^i L_\mathbb{R}/q\Lambda\right)$.

Definition 14. Let Υ be a family of error distributions and let U_Λ be the uniform distribution on $\left(\Lambda_q, \left(\bigoplus_{i=0}^{d-1} u^i L_\mathbb{R}\right)/q\Lambda\right)$. The decision CLWE problem $\mathrm{DCLWE}_{q,s,\psi}$ is, given a number of independent samples from $\Pi_{q,s,\psi}$ for a random pair $(s, \psi) \leftarrow U(\Lambda_q) \times \Upsilon$ or from U_Λ, to decide which with non-negligible advantage.

The hardness of DCLWE was proven in [9] under the canonical embedding[1]. Unlike in that work, here we consider CLWE under the coefficient embedding. This currently lacks a formal security proof, but as explained in the introduction, there is good reason to consider DCLWE a 'hard' problem.

6 The Middle Product for Skew Polynomials

We now define a middle product for skew polynomials. This middle product again takes two (skew) polynomials, multiplies them together, then discards the lower and higher coefficients, forming a (skew) polynomial whose coefficients are the 'middle' part of the product. Below, $\mathcal{R}$ is a ring.

[1] We note here that the reduction required a restriction of the secret space.

Definition 15. Let $d_a, d_b, d, k \in \mathbb{Z}_{\geq 0}$ such that $d_a + d_b - 1 = d + 2k$. The middle-product of $a \in \mathcal{R}^{<d_a}[u,\theta]$ and $b \in \mathcal{R}^{<d_b}[u,\theta]$ is defined

$$\begin{aligned} \odot_d :& \mathcal{R}^{<d_a}[u,\theta] \times \mathcal{R}^{<d_b}[u,\theta] \to \mathcal{R}^{<d}[u,\theta], \\ (a,b) \mapsto a \odot_d b = & \left\lfloor \frac{(a \cdot b) \bmod u^{k+d}}{u^k} \right\rfloor . \end{aligned}$$

We now define skew middle product learning with errors, over $\mathcal{O}_L$:

Definition 16. (SMPLWE distribution) Let $n, d > 0, q \geq 2$, and χ be a distribution over $L_{\mathbb{R}}^{<d}[u,\theta]$. For $s \in \mathcal{O}_{L_q}^{<n+d-1}[u,\theta]$, define the distribution $\mathrm{SMP}_{q,s,n,d,\chi}$ over $\mathcal{O}_{L_q}^{<n}[u,\theta] \times L_{\mathbb{R}_q}^{<d}[u,\theta]$ as the distribution obtained from sampling $a \leftarrow U\left(\mathcal{O}_{L_q}^{<n}[u,\theta]\right)$, $e \leftarrow \chi$ and outputting $(a, b = a \odot_d s + e)$.

Definition 17. (decision SMPLWE) Let $n, d > 0, q \geq 2$, and χ be a distribution over $L_{\mathbb{R}}^{<d}[u,\theta]$. Then decision SMPLWE, $\mathrm{SMPLWE}_{q,s,n,d,\chi}$, consists in distinguishing between arbitrarily many samples from $\mathrm{SMP}_{q,s,n,d,\chi}$ and the same number of samples from $U\left(\mathcal{O}_{L_q}^{<n}[u,\theta] \times L_{\mathbb{R}_q}^{<d}[u,\theta]\right)$, with non-negligible probability over $s \leftarrow U\left(\mathcal{O}_{L_q}^{<n+d-1}[u,\theta]\right)$.

We now prove two lemmas:

Lemma 3. *Let $d, k > 0$, $r \in \mathcal{O}_L^{<k+1}[u,\theta]$, $a \in \mathcal{O}_L^{<k+d}[u,\theta]$, and $b = r \odot_d a$. Let θ be an L-automorphism of order d. We have $\mathbf{b} = \mathrm{GToep}^{d,k+1}(r) \cdot \mathbf{a}$.*

Proof. We can write $r \odot_d a = \sum_{i=0}^{d-1} u^i (\sum_{j+l=i+k} \theta^l(r_j) a_l)$. Thus

$$\begin{aligned} \mathbf{b} = \big(&\theta^k(r_0)a_k + \theta^{k-1}(r_1)a_{k-1} + ... + r_k a_0, \\ &\theta^{k+1}(r_0)a_{k+1} + \theta^k(r_1)a_k + ... + \theta(r_k)a_1, \\ &..., \theta^{k+d-1}(r_0)a_{k+d-1} + \theta^{k+d-2}(r_1)a_{k+d-2} + ... + \theta^{d-1}(r_k)a_{d-1}\big). \end{aligned}$$

and this is precisely $\mathrm{GToep}^{d,k+1}(r) \cdot \mathbf{a}$. and the result follows. □

Lemma 4. *(associativity) Let $d, k, n > 0$. For $r \in \mathcal{O}_L^{<k+1}[u,\theta]$, $a \in \mathcal{O}_L^{<n}[u,\theta]$, and $s \in \mathcal{O}_L^{<n+d+k-1}[u,\theta]$, we have $\theta^{n-1}(r) \odot_d (a \odot_{d+k} s) = (r \cdot a) \odot_d s$.*

Proof. First, observe that the left hand side and right hand side have the same degree. Let the vector of $(r \cdot a) \odot_d s$ be denoted by $\mathbf{u}$, that of $\theta^{n-1}(r) \odot_d (a \odot_{d+k} s)$ by $\mathbf{v}$, and that of $a \odot_{d+k} s$ by $\mathbf{w}$.

For $d, k > 0$, and $r \in \mathcal{O}_L^{<k+1}[u,\theta]$, set $\mathrm{HToep}^{d,k+1}(r)$ to be the $d \times (k+d)$ matrix whose i, jth entry is given by $\theta^{k+d-j}((u^{i-1}r)_{j-1})$, where for polynomial f, $(f)_l$ denotes the lth coefficient of f, indexed from 0. This is the matrix such that $\overline{\mathbf{b}} = \mathrm{HToep}^{d,k+1}(r)\overline{\mathbf{a}}$ for $b = r \odot_d a$. We then have

$$\overline{\mathbf{v}} = \mathrm{HToep}^{d,k+1}(\theta^{n-1}(r)) \cdot \overline{\mathbf{w}} = \mathrm{HToep}^{d,k+1}(\theta^{n-1}(r)) \left(\mathrm{HToep}^{d+k,n}(a) \cdot \overline{\mathbf{s}}\right).$$

Moreover, $\overline{\mathbf{u}} = \mathrm{HToep}^{d,k+n}(r \cdot a) \cdot \overline{\mathbf{s}}$. The result follows from the property $\mathrm{HToep}^{d,k+1}(\theta^{n-1}(r))\,\mathrm{HToep}^{d+k,n}(a) = \mathrm{HToep}^{d,k+n}(r \cdot a)$. □

We can view decision SMPLWE$_{q,d,d,\chi}$ as a structured RLWE variant as follows: given polynomially many samples $(\mathrm{GToep}^{d,d}(a_i), \mathbf{b}_i) \in \mathcal{O}_{L_q}^{d\times(2d-1)} \times L_{\mathbb{R}_q}^{d}$ for uniform $a_i \leftarrow U\left(\mathcal{O}_{L_q}^{<d}[u,\theta]\right)$, decide if the $\mathbf{b}_i$ were sampled uniformly over the domain or have the form $\mathbf{b}_i = \mathrm{GToep}^{d,d}(a_i)\,\mathbf{s} + \mathbf{e}_i$ for some uniform $s \leftarrow U\left(\mathcal{O}_{L_q}^{<2d-1}[u,\theta]\right)$ and $e_i \leftarrow \chi$. Note the samples are correlated.

7 Reduction from SPLWE to SMPLWE

We adapt the reduction for standard MPLWE, under the coefficient embedding.

Theorem 1. *Let $d > 0, q \geq 2$, and χ a distribution over $L_{\mathbb{R}}^{<d}[u,\theta]$. Then there exists a ppt. reduction from SPLWE$_{q,s,f,\chi}$ for any polynomial of the form $f(u) = u^d - \gamma \in \mathcal{O}_L[u,\theta]$ with $\gamma \in \mathcal{O}_K \setminus \{0\}$ coprime with q, to SMPLWE$_{q,s,d,d,\chi'}$.*

Proof. Like in [19], we use an efficiently computable transformation ϕ that maps $(a_i, b_i) \in \mathcal{O}_{L_q}[u,\theta]/f \times L_{\mathbb{R}_q}[u,\theta]/f$ to $(a_i', b_i') \in \mathcal{O}_{L_q}^{<d}[u,\theta] \times L_{\mathbb{R}_q}^{<d}[u,\theta]$, sending $U\left(\mathcal{O}_{L_q}[u,\theta]/f \times L_{\mathbb{R}_q}[u,\theta]/f\right)$ to $U(\mathcal{O}_{L_q}^{<d}[u,\theta] \times L_{\mathbb{R}_q}^{<d}[u,\theta])$ and SP$_{q,s,f,\chi}$ to SMP$_{q,s',d,d,\chi'}$, for a new s' that is a function of s and a new distribution χ' that depends on χ and f. Given such a ϕ, the steps of the reduction are:

1. Sample a uniform $t \leftarrow U\left(\mathcal{O}_{L_q}^{<2d-1}[u,\theta]\right)$.
2. For each SPLWE sample (a_i, b_i), compute $(a_i, b_i') = \phi(a_i, b_i)$. Give $(a_i, b_i') + (0, \tilde{a_i} \odot_d t)$ to the SMPLWE oracle.
3. Return the output of the oracle.

For such a transformation ϕ, the reduction preserves the uniformity of uniform samples, and maps SP$_{q,s,f,\chi}$ samples to SMP$_{q,s'+t,d,d,M_{f,\theta}\cdot\chi}$ samples. When s is uniform, the SMP$_{q,s'+t,d,d,M_{f,\theta}\cdot\chi}$ samples have a uniform $s' + t$.

To construct ϕ, let $(a_i, b_i) \in \mathcal{O}_{L_q}[u,\theta]/f \times L_{\mathbb{R}_q}[u,\theta]/f$ be a SPLWE sample. Let $\deg(f) = d$. Set $\phi(a_i, b_i) = (a_i, b_i')$ where b_i' is defined

$$\mathbf{b}_i' = M_{f,\theta} \cdot \mathbf{b}_i \in L_{\mathbb{R}_q}^{<d}[u,\theta].$$

Plainly a_i is uniform, by definition. Observe that if b_i is uniformly distributed, then so is its vector of coefficients $\mathbf{b}_i$. Moreover, since the matrix $M_{f,\theta}$ is invertible modulo q we find $M_{f,\theta} \cdot \mathbf{b}_i$ is also uniform.

Now write $b_i = a_i \cdot s + e_i$, for $s \in \mathcal{O}_{L_q}[u,\theta]/f$ and $e_i \leftarrow \chi$. Since $\mathrm{Rot}_f(b_i) = \mathrm{Rot}_f(a_i) \cdot \mathrm{Rot}_f(s) + \mathrm{Rot}_f(e_i)$, we have

$$\begin{aligned}
M_{f,\theta} \cdot \mathbf{b}_i &= \mathrm{Rot}_f(b_i) \cdot (1, 0, .., 0)^T \\
&= (\mathrm{Rot}_f(a_i) \cdot \mathrm{Rot}_f(s) + \mathrm{Rot}_f(e_i)) \cdot (1, 0, ..., 0)^T \\
&= \mathrm{Rot}_f(a_i) \cdot \mathrm{Rot}_f(s) \cdot (1, 0, ..., 0)^T + \mathrm{Rot}_f(e_i) \cdot (1, 0, ..., 0)^T \\
&= \mathrm{Rot}_f(a_i) \cdot M_{f,\theta} \cdot \mathbf{s} + M_{f,\theta} \cdot \mathbf{e}_i \\
&= \mathrm{GToep}^{d,d}(\tilde{a_i}) \cdot N_f \cdot M_{f,\theta} \cdot \mathbf{s} + M_{f,\theta} \cdot \mathbf{e}_i \\
&= \mathrm{GToep}^{d,d}(\tilde{a_i}) \cdot \mathbf{s}' + M_{f,\theta} \cdot \mathbf{e}_i,
\end{aligned}$$

where $\mathbf{s}' = N_f \cdot M_{f,\theta} \cdot \mathbf{s}$. Since $\mathbf{b}_i' = M_{f,\theta} \cdot \mathbf{b}_i = \mathrm{GToep}^{d,d}\left(\tilde{a_i}\right) \cdot \mathbf{s}' + M_{f,\theta} \cdot \mathbf{e}_i$, the new error is $\mathbf{e}_i' = M_{f,\theta} \cdot \mathbf{e}_i$, as required □

In order to remove dependence on the choice of γ, one can consider a family of polynomials $\mathcal{F}_\beta := \{f(u) = u^d - \gamma : |\gamma| \leq \beta\}$. If $\chi = D_{\alpha q}$, then $\chi' = M_{f,\theta} \cdot D_{\alpha q}$. Expanding $M_{f,\theta}$ over $\mathbb{Z}$, since $M_{f,\theta}$ is invertible, we have $\chi' = D_{M_{f,\theta} \cdot (\alpha q I_{[L:\mathbb{Q}]})}$. Since the the square of the largest singular value $\|M_{f,\theta}\|^2 = |\gamma|^2$, then restricting to $f \in \mathcal{F}_\beta$, adding an error $e_i' \leftarrow D_\Sigma$ for a positive definite Σ such that $M_{f,\theta} \cdot \boldsymbol{e_i} + e_i' \sim D_{\alpha q \beta}$ removes any dependence of the error on the choice of $f \in \mathcal{F}_\beta$.

8 Public Key Encryption Scheme

In this section we give an encryption scheme and prove its IND-CPA security. Let L/K be a cyclic Galois extension of degree d, $\mathrm{Gal}(L/K) = \langle\theta\rangle$, $[K : \mathbb{Q}] = n$, and q unramified in $\mathcal{O}_L$. The scheme uses the following error distribution: let $\chi = \lfloor D_{\alpha q} \rceil$ be a discretised Gaussian over $\mathcal{O}_L^{<d+k}[u,\theta]$, where coefficients are sampled from $D_{\alpha q}$, rounded to the nearest integer, and set as the $\mathbb{Z}$-coefficients of a skew polynomial in $\mathcal{O}_L^{<d+k}[u,\theta]$. Plaintexts are taken from $\mathcal{B}^{<d}[u,\theta]$, where $\mathcal{B} = \{a(x) \in \mathcal{O}_L : a_i \in \{0,1\}$ for all $i\}$. We denote $\mathcal{B}^\times := \mathcal{B} \bmod q\mathcal{O}_L \cap \mathcal{O}_{L_q}^\times$. Ciphertexts will be elements of $\mathcal{O}_{L_q}^{<d+2k}[u,\theta] \times \mathcal{O}_{L_q}^{<d}[u,\theta]$.

Key Generation To generate a key pair (pk, sk), begin by sampling $s \leftarrow U\left(\mathcal{O}_{L_q}^{<2(d+k)-1}[u,\theta]\right)$. Then for all $i \leq t$, sample uniform $a_i \leftarrow U\left(\mathcal{O}_{L_q}^{<d+k}[u,\theta]\right)$ and errors $e_i \leftarrow \chi$, and set $b_i = a_i \odot_{d+k} s + 2 \cdot e_i \in \mathcal{O}_{L_q}^{<d+k}[u,\theta]$, $i = 1, ..., t$. The public key is $pk := (a_i, b_i)_{i \leq t}$, and the secret key is $sk := s$.

Encryption Given public key $pk = (a_i, b_i)_{i \leq t}$ we encrypt a message $\mu \in \mathcal{B}^{<d}[u,\theta]$ as follows. We sample $r_i \leftarrow U\left(\mathcal{B}^{<k+1}[u,\theta]\right)$, $i = 1, ..., t$, replace the smallest non-zero $\mathcal{O}_L$-coefficient of each r_i with an element sampled uniformly from $\mathcal{B}^\times$, and output a ciphertext $c = (c_1, c_2) \in \mathcal{O}_{L_q}^{<d+2k}[u,\theta] \times \mathcal{O}_{L_q}^{<d}[u,\theta]$, where

$$c_1 = \sum_{i \leq t} r_i \cdot a_i, \text{ and } c_2 = \mu + \sum_{i \leq t} \theta^{d+k-1}(r_i) \odot_d b_i$$

Decryption Given $sk = s$, to decrypt a ciphertext $c = (c_1, c_2)$, compute

$$\mu' := (c_2 - c_1 \odot_d s \bmod q) \bmod 2$$

We now show correctness.

Lemma 5. *Let $\alpha < 1/(16\sqrt{ndt(k+1)})$ and $q \geq 16ndt(k+1)$. With probability at least $1 - nd^2 \cdot 2^{-\Omega(n)}$ over valid key pairs (pk, sk), for all plaintexts $\mu \in \mathcal{B}^{<d}[u,\theta]$ and with probability 1 over the encryption randomness, decryption is correct.*

Proof. Suppose that $c = (c_1, c_2)$ is a ciphertext encrypting a message μ under a public key $pk = (a_i, b_i)_{i \leq t}$. Then to decrypt c we compute

$$\begin{aligned} c_2 - c_1 \odot_d s &= \mu + \sum_{i \leq t} \theta^{d+k-1}(r_i) \odot_d b_i - \left(\sum_{i \leq t} r_i \cdot a_i \right) \odot_d s \\ &= \mu + \sum_{i \leq t} \left(\theta^{d+k-1}(r_i) \odot_d (a_i \odot_{d+k} s + 2 \cdot e_i) - (r_i \cdot a_i) \odot_d s \right) \\ &= \mu + \sum_{i \leq t} \theta^{d+k-1}(r_i) \odot_d (a_i \odot_{d+k} s) - (r_i \cdot a_i) \odot_d s + 2\theta^{d+k-1}(r_i) \odot_d e_i \\ &= \mu + 2 \sum_{i \leq t} \theta^{d+k-1}(r_i) \odot_d e_i \bmod q \end{aligned}$$

where the final equality holds by Lemma 4. Note that if

$$\|\mu + 2 \cdot \sum_{i \leq t} \theta^{d+k-1}(r_i) \odot_d e_i\|_\infty < q/2,$$

then $c_2 - c_1 \odot_d s \bmod q = \mu + 2 \cdot \sum_{i \leq t} \theta^{d+k-1}(r_i) \odot_d e_i$, so $c_2 - c_1 \odot_d s \bmod q \bmod 2 = \mu$. Similarly to [19, Lemma 4.1], the coefficients of $\sum_{i \leq t} \theta^{d+k-1}(r_i) \odot_d e_i$ can be written as an inner product between a binary $[\mathcal{O}_L : \mathbb{Z}]t(k+1)$-dimensional vector and a vector distributed according to $\lfloor D_{\alpha q} \rceil^{[\mathcal{O}_L : \mathbb{Z}]t(k+1)}$, so applying a (Gaussian) tail bound and the triangle inequality, the coefficients each have magnitude at most $\alpha q \sqrt{[\mathcal{O}_L : \mathbb{Z}]t(k+1)} + [\mathcal{O}_L : \mathbb{Z}]t(k+1)$ with probability at least $1 - 2^{-\Omega(n)}$. Thus $\|\mu + 2 \cdot \sum_{i \leq t} \theta^{d+k-1}(r_i) \odot_d e_i\|_\infty < 2\alpha q \sqrt{[\mathcal{O}_L : \mathbb{Z}]t(k+1)} + 2t[\mathcal{O}_L : \mathbb{Z}](k+1) + 1$ with probability at least $1 - d[\mathcal{O}_L : \mathbb{Z}]2^{-\Omega(n)}$. □

To show security of the above scheme, we demonstrate its IND-CPA security, assuming the hardness of SMPLWE, following [19]. We denote the set of r_i obtainable during the encryption procedure by $\overline{\mathcal{B}}^{<k+1}[u, \theta]$, and write $r_i \leftarrow \overline{\mathcal{B}}^{<k+1}[u, \theta]$.

Lemma 6. *Let $q, k, d \geq 2$. For $b_i \in \mathcal{O}_{L_q}^{<d+k}[u, \theta]$, let h_{b_i} denote the map that sends $r_i \leftarrow \overline{\mathcal{B}}^{<k+1}[u, \theta]$ to $r_i \odot_d b_i \in \mathcal{O}_{L_q}^{<d}[u, \theta]$. Then the hash function family $\mathcal{H} = (h_{b_i})_{b_i}$ is universal.*

Proof. Identical to [19, Lemma 4.2], included for completeness. It suffices to prove that for all $y \in \mathcal{O}_{L_q}^{<d}[u, \theta]$

$$\Pr_{b_1} [r_1 \odot_d b_1 = y] = |\mathcal{O}_{L_q}|^{-d}.$$

Let j be the smallest integer such that the u^j-coefficient of r_1 is non-zero and let r_1 have ith coefficient $r_{1,i}$. Then $r_1 \odot_d b_1 = y$ restricted to entries $j, ..., j+d-1$ can be written as a triangular linear map with entries in $\{r_{1,j}, ..., r_{1,j+d-1}\}$ and $r_{1,j}$ along the diagonal, applied to the vector of d coefficients of b_1, up to application of θ. Since $r_{1,j}$ is invertible by construction, restricting the map $b_1 \mapsto r_1 \odot_d b_1$ to these d coefficients of b_1 is a bijection, which implies the result. □

By linearity the hash function family $(h_{(b_i)_i})_{(b_i)_i}$ with $(b_i)_i \in \left(\mathcal{O}_{L_q}^{<d+k}[u,\theta]\right)^t$ and h_{b_i} mapping $(r_i)_{i\leq t} \leftarrow \left(\overline{\mathcal{B}}^{<k+1}[u,\theta]\right)^t$ to $\sum_i r_i \odot_d b_i$ is also universal.

Theorem 2. *Let $t \geq \frac{2+2(k+d)\log(q)}{k}$. Then the SMPLWE PKE scheme is IND-CPA secure, assuming the hardness of SMPLWE$_{q,d+k,d+k,D_{\alpha q}}$.*

Proof. We perform two hops from the IND-CPA experiment for SMPLWE to an experiment which we show to be of negligible statistical distance from our starting point. We first consider a variant of the IND-CPA experiment in which $pk = (a_i, b_i)_i$ is sampled uniformly. Assuming the hardness of decision SMPLWE, the probabilities that $\mathcal{A}$ outputs $b' = b$ in the IND-CPA experiment and in the variant experiment are negligibly close.

Now consider a second experiment. Suppose $pk = (a_i, b_i)_i$ is a valid public key, but instead of computing a valid ciphertext c encrypting μ_b under pk for $b \in \{0,1\}$, $c = (c_1, c_2)$ is computed by the following process: sample uniform $r_i \leftarrow \overline{\mathcal{B}}^{<k+1}[u,\theta]$, $i = 1, ..., t$, sample a uniform $v \leftarrow U\left(\mathcal{O}_{L_q}^{<d}[u,\theta]\right)$, and set

$$(c_1, c_2) := \left(\sum_{i=1}^{t} r_i \cdot a_i, v\right)$$

Since v is independent of b, the probability that $\mathcal{A}$ outputs $b' = b$ is precisely 1/2. We now show that the distributions of $((a_i, b_i)_i, c_1, c_2)$ in the two variant experiments are of negligible statistical distance from one another; that is, that

$$\Delta\left(\left((a_i, b_i)_i, \sum_{i\leq t} r_i \cdot a_i, \sum_{i\leq t} r_i \odot_d b_i\right), \quad \left((a_i, b_i)_i, \sum_{i\leq t} r_i \cdot a_i, v\right)\right) \leq \mathrm{neg}(n)$$

where a_i, b_i, r_i, and v are sampled uniformly from $\mathcal{O}_{L_q}^{<d+k}[u,\theta]$, $\mathcal{O}_{L_q}^{<d+k}[u,\theta]$, $\overline{\mathcal{B}}^{<k+1}[u,\theta]$ and $\mathcal{O}_{L_q}^{<d}[u,\theta]$ respectively, for $i = 1, ..., t$. Applying Lemma 1, since Lemma 6 showed the hash function family $(h_{b_i})_{b_i}$ is universal, and noting that $\sum_{i\leq t} r_i \cdot a_i \in \mathcal{O}_{L_q}^{<d+2k}[u,\theta]$ which is of cardinality $|\mathcal{O}_{L_q}|^{d+2k}$, we find that the statistical distance above is upper bounded by $\frac{1}{2}\sqrt{\gamma(T) \cdot |Y| \cdot |Z|}$, where $X = (\overline{\mathcal{B}}^{<k+1}[u,\theta])^t$, $\gamma(T) = \max_{w\in X} \Pr[T = w] \leq |\mathcal{B}|^{-tk}$, $|Y| = |\mathcal{O}_{L_q}|^d$, and $|Z| = |\mathcal{O}_{L_q}|^{d+2k}$; so the upper bound is

$$\frac{1}{2}\left(|\mathcal{B}|^{-tk} \cdot |\mathcal{O}_{L_q}|^{2(d+k)}\right)^{1/2} = \frac{1}{2}\left(2^{-ndtk} \cdot q^{2nd(d+k)}\right)^{1/2}$$

If $t \geq \frac{2+2(k+d)\log(q)}{k}$ this becomes negligible in n. □

9 Conclusion

We have introduced SMPLWE and SPLWE and reduced a family of problems based on the latter to the former. We have connected SPLWE and CLWE. We also gave a PKE scheme and proved its security under a reasonable assumption. Future work might include removing restrictions on the degrees of the polynomials involved, and obtaining greater functionality from the SMPLWE problem.

References

1. Ajtai, M.: Generating hard instances of lattice problems. Electron. Colloquium Comput. Complex. **TR96** (1996). https://doi.org/10.1145/237814.237838
2. Albert, A.: Structure of Algebras, AMS colloquium publications, vol. 24. American Mathematical Society (1939)
3. Augot, D., Loidreau, P., Robert, G.: Generalized gabidulin codes over fields of any characteristic. Des. Codes Cryptography **86**(8), 1807–1848 (2018). https://doi.org/10.1007/s10623-017-0425-6
4. Bai, S., Boudgoust, K., Das, D., Roux-Langlois, A., Wen, W., Zhang, Z.: Middle-product learning with rounding problem and its applications. In: Galbraith, S., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11921, pp. 55–81. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-34578-5$_3$
5. Bai, S., Das, D., Hiromasa, R., Roşca, M., Sakzad, A., Stehlé, D., Steinfeld, R., Zhang, Z.: MPSign: A signature from small-secret middle-product learning with errors. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) PKC 2020. LNCS, vol. 12111, pp. 66–93. Springer International Publishing (2020)
6. Blanco-Chacón, I.: On the RLWE/PLWE equivalence for cyclotomic number fields. Applicable Algebra in Engineering, Communication and Computing **33** (01 2022). https://doi.org/10.1007/s00200-020-00433-z
7. Caruso, X., Le Borgne, J.: Fast multiplication for skew polynomials. In: Burr, M., Yap, C.K., Din, M.S.E. (eds.) ISSAC 2017. p. 77–84. Association for Computing Machinery (2017). https://doi.org/10.1145/3087604.3087617
8. Giorgi, P.: A probabilistic algorithm for verifying polynomial middle product in linear time. Information Processing Letters **139**, 30–34 (2018). https://doi.org/10.1016/j.ipl.2018.06.014
9. Grover, C., Mendelsohn, A., Ling, C., Vehkalahti, R.: Non-commutative ring learning with errors from cyclic algebras. J. of Cryptology **35**(3), 22 (2022). https://doi.org/10.1007/s00145-022-09430-6
10. Hanrot, G., Quercia, M., Zimmermann, P.: The middle product algorithm I. Appl. Algebra Eng., Commun. Comput. **14**(6), 415–438 (2004). https://doi.org/10.1007/s00200-003-0144-2
11. Huffman, W., Kim, J., Solé, P.: Concise Encyclopedia of Coding Theory. CRC Press (2021)
12. Katz, J., Lindell, Y.: Introduction to Modern Cryptography, 2nd Edition. Chapman & Hall/CRC Cryptography and Network Security Series, Taylor & Francis (2014)
13. Ling, C., Mendelsohn, A.: NTRU in quaternion algebras of bounded discriminant. In: Johansson, T., Smith-Tone, D. (eds.) PQCrypto 2023. LNCS, vol. 14154, pp. 256–290. Springer Nature Switzerland (2023)
14. Lombardi, A., Vaikuntanathan, V., Vuong, T.: Lattice trapdoors and IBE from middle-product LWE. In: Hofheinz, D., Rosen, A. (eds.) TCC 2019. LNCS, vol. 11891, pp. 24–54. Springer International Publishing (2019)
15. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer Berlin Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5$_1$

16. Ore, O.: Theory of non-commutative polynomials. Annals of Mathematics **34**(3), 480–508 (1933), `http://www.jstor.org/stable/1968173`
17. Peikert, C., Pepin, Z.: Algebraically structured LWE, revisited. In: Hofheinz, D., Rosen, A. (eds.) TCC 2019. LNCS, vol. 11891, pp. 1–23. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-36030-6$_1$
18. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. Journal of the ACM **56** (2009). https://doi.org/10.1145/1568318
19. Roșca, M., Sakzad, A., Stehlé, D., Steinfeld, R.: Middle-product learning with errors. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10403, pp. 283–297. Springer International Publishing (2017)
20. Roșca, M., Stehlé, D., Wallet, A.: On the Ring-LWE and Polynomial-LWE problems. In: Nielsen, J., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10820, pp. 146–173. Springer International Publishing (2018)
21. Shor, P.W.: Algorithms for quantum computation: discrete logarithms and factoring. In: 35th FOCS. pp. 124–134. IEEE Computer Society Press (Nov 1994)
22. Stehlé, D., Steinfeld, R., Tanaka, K., Xagawa, K.: Efficient public key encryption based on ideal lattices. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 617–635. Springer Berlin Heidelberg (2009)
23. Steinfeld, R., Sakzad, A., Zhao, R.K.: Practical MP-LWE-based encryption balancing security-risk versus efficiency. Des. Codes Cryptogr. **87**, 2847–2884 (2019). https://doi.org/10.1007/s10623-019-00654-5

A Skew Polynomial Rings

In this appendix we give a fuller explanation of the theory of skew polynomials.

Definition 18. Let R be a commutative ring. A polynomial in the indeterminate x with coefficients in R is an expression of the form

$$a_0 + a_1x + ... + a_nx^n,$$

where x commutes with elements of R, $a_i \in R$ for $i = 0, ..., n$, and $n < \infty$.

We call n the degree of the polynomial, and if we label $f(x) = a_0+a_1x+...+a_nx^n$, then we write $\deg(f) = n$. The set of polynomials with coefficients in R is denoted $R[x]$. This set has a ring structure, where addition is performed coefficient-wise (e.g. $a_0 + a_1x + b_0 + b_1x = a_0 + b_0 + (a_1 + b_1)x$) and multiplication is defined

$$(a_0 + a_1x + ... + a_nx^n) \cdot (b_0 + b_1x + ... + b_mx^m) = \sum_{k=0}^{n+m} \sum_{l=0}^{k} a_lb_{k-l}x^k$$

Definition 19. If R and S are two rings, we let $\mathrm{Hom}(R, S)$ denote the set of homomorphisms from R to S and $\mathrm{Iso}(R, S)$ the set of isomorphisms from R to S. If $R = S$, then we write $\mathrm{End}(R) = \mathrm{Hom}(R, R)$ for the endomorphisms of R and $\mathrm{Aut}(R) = \mathrm{Iso}(R, R)$ for the automorphisms of R.

Let $\mathbb{F}'$ be an algebraic field extension of $\mathbb{F}$. Then any $\mathbb{F}$-endomorphism of $\mathbb{F}'$ is an $\mathbb{F}$-automorphism of $\mathbb{F}'$.

The order of an endomorphism θ is the smallest integer d such that $\theta^d = \mathrm{id}$.

Examples 1. Let $\mathbb{C}$ denote the complex numbers and $\bar{\cdot}$ complex conjugation, that is, the map sending $a + ib \mapsto a - ib =: \overline{a + ib}$. Then $\bar{\cdot}$ is an automorphism of $\mathbb{C}$, and has order 2.
2. Let $\mathbb{F}_q$ be a finite extension of $\mathbb{F}_p$, the finite field of p elements. Then the map $a \mapsto a^p$ is an automorphism of $\mathbb{F}_q$, called the Frobenius map, denoted Frob_p. If $q = p^r$, Frob_p has order r.
3. Let $\mathbb{Q}(\sqrt{d})$ be a real quadratic extension of $\mathbb{Q}$ with defining polynomial $f(x) = x^2 - d$ for some $d \in \mathbb{N}$. Then the map τ sending $d \mapsto -d$ and fixing $\mathbb{Q}$ is an automorphism of $\mathbb{Q}(\sqrt{d})$ of order 2.

Definition 20. Let R be a ring and θ an endomorphism of R. Then expressions in the indeterminate x of the form

$$a_0 + a_1 x + ... + a_n x^n$$

where $xr = \theta(r)x$ for all $r \in R$, $a_i \in R$ for $i = 0, ..., n$, and $n < \infty$ are called skew polynomials.

The degree of a skew polynomial $f(x) = a_0 + a_1 x + ... + a_n x^n$ is n. We denote the set of skew polynomials with coefficients in R and indeterminate x defined by some endomorphism θ by $R[x, \theta]$. If θ is the identity map id, then $R[x, \text{id}] = R[x]$.

Proposition 3. *Let R be a ring and $\theta \in \text{End}(R)$. Then $R[x, \theta]$ is a ring.*

Proof. Addition is coefficient-wise (e.g. $a_0 + a_1 x + b_0 + b_1 x = a_0 + b_0 + (a_1 + b_1)x$). Multiplication is defined

$$(a_0 + a_1 x + ... + a_n x^n) \cdot (b_0 + b_1 x + ... + b_m x^m) = \sum_{k=0}^{n+m} \sum_{l=0}^{k} a_l \theta^l \left(b_{k-l}\right) x^k$$

The result follows from axiom checking. □

Let R be an integral domain and θ be injective. Then $a_n \theta^n(b_m) \neq 0$ if $a_n, b_m \neq 0$, so the leading term of the product of $a_0 + a_1 x + ... + a_n x^n$ and $b_0 + b_1 x + ... + b_m x^m$ is non-zero. This allows us to generalise the notion of degree to skew polynomials. Thus when R is a domain and θ injective the degree of the above product is $n+m$ and the degree of the product of two skew polynomials is the sum of the degrees.

Examples 1. $\mathbb{C}[x, \bar{\cdot}]$. Write $\iota(\cdot) = \bar{\cdot}$ for convenience. We have

$$\begin{aligned}(a_0 + ... + a_n x^n) \cdot (b_0 + ... + b_m x^m) &= \sum_{k=0}^{n+m} \sum_{l=0}^{k} a_l \iota^l (b_{k-l}) x^k \\ &= \sum_{k=0}^{n+m} \left(\sum_{l \text{ even}} a_l b_{k-l} + \sum_{l \text{ odd}} a_l \overline{b_{k-l}} \right) x^k\end{aligned}$$

2. $\mathbb{F}_{p^r}[x, \mathrm{Frob}_p]$. Then $(a_0 + ... + a_n x^n) \cdot (b_0 + ... + b_m x^m) = \sum_{k=0}^{n+m} \sum_{l=0}^{k} a_l b_{k-l}^{p^l} x^k$.
3. $\mathbb{Q}(\sqrt{d})[x, \tau]$. Then

$$\begin{aligned}(a_0 + ... + a_n x^n) \cdot (b_0 + ... + b_m x^m) &= \sum_{k=0}^{n+m} \sum_{l=0}^{k} a_l \tau^l (b_{k-l}) x^k \\ &= \sum_{k=0}^{n+m} \left(\sum_{l \text{ even}} a_l b_{k-l} + \sum_{l \text{ odd}} a_l \tau(b_{k-l}) \right) x^k\end{aligned}$$

A left ideal $\mathcal{I}$ of a ring R is an additively closed subgroup which is closed under multipliction on the left from R, that is, $R\mathcal{I} \subset \mathcal{I}$. Right ideals are defined analogously. An ideal is principal if it is generated by a single element. We have

Proposition 4. *If R is an integral domain and θ is injective, then $R[x, \theta]$ is an integral domain. If K is a field and σ an endomorphism of K, then every left ideal of $K[x, \sigma]$ is principally generated.*

The above gives an analogous statement to the fact that a polynomial ring $K[x]$ over a (commutative) field K is a PID. A similar statement holds for right ideals.

Definition 21. Let R be a ring and $\theta \in \mathrm{End}(R)$. Then we call

$$R^\theta := \{y \in R : \theta(y) = y\}$$

the fixed ring of θ.

Note the above is a ring: $0, 1 \in R^\theta$, R^θ inherits associativity and distributivity from R, and is additively and multiplicatively closed by the properties of θ. If K is a field and $\sigma \in \mathrm{Aut}(K)$, K^σ is a subfield of K called the fixed field of σ.

Definition 22. The center $\mathcal{Z}(R)$ of a (noncommutative) ring R is the set of elements of R which commute with all other elements of R; that is,

$$\mathcal{Z}(R) := \{y \in R : yz = zy \text{ for all } z \in R\}$$

It is clear that $\mathcal{Z}(R)$ is a commutative subring of R. The following describes the center of a skew polynomial ring:

Proposition 5. *Let R be a ring and $\theta \in \mathrm{End}(R)$ have finite order d. Then the center of $R[x, \theta]$ is given by $\mathcal{Z}(R[x, \theta]) = \mathcal{Z}(R)[x^d]$. If θ has infinite order, then $\mathcal{Z}(R[x, \theta]) = \mathcal{Z}(R)$.*

A central element z generates a two-sided ideal, since $Rz = zR$ by definition.

Examples 1. $\mathcal{Z}(\mathbb{C}[x, \bar{\cdot}])$. The fixed field of $\bar{\cdot}$ is $\mathbb{R}$, since $\overline{a + i \cdot 0} = \overline{a} = a$. Since $\bar{\cdot}$ has order two, we find $\mathcal{Z}(\mathbb{C}[x, \bar{\cdot}]) = \mathbb{R}[x^2]$.
2. $\mathcal{Z}(\mathbb{F}_{p^r}[x, \mathrm{Frob}_p])$. The fixed field of Frob_p is $\mathbb{F}_p$ and Frob_p has order r, so we find $\mathcal{Z}(\mathbb{F}_{p^r}[x, \mathrm{Frob}_p]) = \mathbb{F}_p[x^r]$.
3. $\mathcal{Z}(\mathbb{Q}(\sqrt{d})[x, \tau])$. Since $\mathbb{Q}(\sqrt{d})^\tau = \mathbb{Q}$ and $\tau^2 = \mathrm{id}$, $\mathcal{Z}(\mathbb{Q}(\sqrt{d})[x, \tau]) = \mathbb{Q}[x^2]$.

We briefly consider some further properties of skew polynomial rings. We first note that Hilbert's basis theorem holds:

Theorem 3. *Let R be a Noetherian ring, θ an automorphism of R, and $S = R[x, \theta]$. Then S is Noetherian.*

Let K be an algebraic number field Galois over $\mathbb{Q}$ and $\mathcal{O}_K$ the ring of integers of K. The $\mathbb{Q}$-automorphisms of K restrict to endomorphisms of $\mathcal{O}_K$, so we can consider the skew polynomial ring $\mathcal{O}_K[x, \theta]$ where $\theta \in \mathrm{Gal}(K/\mathbb{Q})$. Since $\mathcal{O}_K$ is Noetherian, by the theorem so is $\mathcal{O}_K[x, \theta]$.

Let $f, g \in R[x, \theta]$. We say g is a left divisor of f if $f = gh$ for some $h \in R[x, \theta]$. A skew polynomial f is irreducible if all its left divisors are either units or skew polynomials of the same degree as f. Then

Theorem 4. *[16] Let $f_1, ..., f_n$, $g_1, ..., g_m$ be irreducible skew polynomials such that $f_1 \cdot ... \cdot f_n = g_1 \cdot ... \cdot g_m$. Then $n = m$ and $\deg(f_i) = \deg(g_{\pi(i)})$ for some permutation π and $i = 1, ..., n$.*

We can consider quotients of skew polynomial rings. If $\mathcal{I}$ is a left ideal of $R[x, \theta]$, then $R[x, \theta]/\mathcal{I}$ is a left $R[x, \theta]$-module, since if $f(x), g(x) \in R[x, \theta]$

$$f(x)(g(x) + \mathcal{I}) = f(x)g(x) + f(x)\mathcal{I} \subset f(x)g(x) + \mathcal{I}$$

If $\mathcal{I}$ is a two-sided ideal, then $R[x, \theta]/\mathcal{I}$ is a ring:

$$(f(x) + \mathcal{I})(g(x) + \mathcal{I}) = f(x)g(x) + \mathcal{I}g(x) + f(x)\mathcal{I} + \mathcal{I}^2 \subset f(x)g(x) + \mathcal{I}$$

When K is a field, every ideal is principally generated, and so if $z \in \mathcal{Z}(K[x, \sigma])$, then $K[x, \sigma]/zK[x, \sigma]$ is a ring.

Examples 1. Note that $x^2 + \pi \in \mathcal{Z}(\mathbb{C}[x, \bar{\cdot}])$, so $\mathbb{C}[x, \bar{\cdot}]/(x^2 + \pi)\mathbb{C}[x, \bar{\cdot}]$ is a ring.
2. Since $x^{r^2} + 1 \in \mathcal{Z}(\mathbb{F}_{p^r}[x, \mathrm{Frob}_p])$, $\mathbb{F}_{p^r}[x, \mathrm{Frob}_p]/(x^{r^2} + 1)\mathbb{F}_{p^r}[x, \mathrm{Frob}_p]$ is a ring.
3. Note that $x^8 + 1 \in \mathcal{Z}(\mathbb{Q}(\sqrt{d})[x, \tau])$, so $\mathbb{Q}(\sqrt{d})[x, \tau]/(x^8 + 1)\mathbb{Q}(\sqrt{d})[x, \tau]$ is a ring.

B On the Equivalence of Embeddings for CLWE

In [20], [6] instances of number fields were given for which the distortion induced by mapping between the canonical and the coefficient embedding was polynomially bounded, impyling a polynomial-time equivalence between solving RLWE in those fields and solving the corresponding PLWE instances. They achieved this by bounding Frobenius norm of the map V_f which sends the canonical embedding of an element x to a coefficient representation of x, that is

$$\sigma_L(x) = V_f \cdot \mathrm{coeff}(x),$$

where $\mathrm{coeff}(\cdot)$ is the vector of coefficients of $x \in \mathbb{Z}[x]/f(x)$ and σ_L is the canonical embedding. In this appendix, we give examples of CDAs for which the coefficient representation of an algebra element is only polynomially distorted by mapping it into canonical space. These instances were studied in [13].

In particular, we consider CDAs obtained from quadratic extensions of power-of-two conductor cyclotomic fields $K = \mathbb{Q}(\zeta_{2^r})$, obtained by adjoining $\sqrt{\ell}$ to K, where $\ell > 2$ is prime and satisfies $\ell \equiv 1 \bmod 2^r$, $\ell \not\equiv 1 \bmod 2^{r+1}$. Then $\mathcal{A} = (L/K, \theta, \zeta_n)$ is a CDA and Λ is a maximal order in $\mathcal{A}$, with $L = \mathbb{Q}(\zeta_{2^r}, \sqrt{\ell})$.

Write $m = 2^r$ and $n = 2^{r-1}$. We then define the powerful basis of $\mathcal{O}_L$:

$$\overrightarrow{p} := (1, \zeta_m, ..., \zeta_m^{n-1}, \frac{1+\sqrt{\ell}}{2}, \zeta_m \frac{1+\sqrt{\ell}}{2}, ..., \zeta_m^{n-1} \frac{1+\sqrt{\ell}}{2})$$

From this we obtain a matrix in $\mathbb{R}^{n \times n}$ by applying the canonical embedding to the entries of $\overrightarrow{p}$:

$$\sigma_L(\overrightarrow{p}) = \left(\sigma_L(1), ..., \sigma_L(\zeta_m^{n-1}), \sigma_L\left(\frac{1+\sqrt{\ell}}{2}\right), ..., \sigma_L\left(\zeta_m^{n-1}\frac{1+\sqrt{\ell}}{2}\right)\right)$$

It can be checked that $\sigma_L(x) = \sigma_L(\overrightarrow{p}) \cdot \mathrm{coeff}(x)$. This implies that $\|\sigma_L(x)\| \leq s_1(\sigma_L(\overrightarrow{p})) \cdot \|x\|_{\overrightarrow{p}}$, where $\|\cdot\|_{\overrightarrow{p}}$ denotes taking the ℓ_2-norm of the coefficient vector of an element expressed in the basis $\overrightarrow{p}$, and $s_1(\sigma_L(\overrightarrow{p}))$ is the largest singular value of $\sigma_L(\overrightarrow{p})$. Labelling the smallest singular value by $s_{2n}(\cdot)$, we have

Proposition 6. *[13, Proposition 1] Let $n = 2^{r-1}$, $\ell \equiv 1 \bmod 2^r$ a prime, and $L = \mathbb{Q}(\zeta_{2^r}, \sqrt{\ell})$. Then, using the powerful basis of $\mathcal{O}_L$, we have*

$$s_1(\overrightarrow{p}) = \frac{\sqrt{n}}{2}\sqrt{\ell + 5 + \sqrt{\ell^2 - 6\ell + 25}},$$

$$s_{2n}(\overrightarrow{p}) = \frac{\sqrt{n}}{2}\sqrt{\ell + 5 - \sqrt{\ell^2 - 6\ell + 25}}.$$

Therefore for bounded values of ℓ, say $\ell = \mathrm{poly}(n)$, the singular values are also polynomial in n. Bounding the $s_i(\sigma_L(\overrightarrow{p}))$ allows us to bound V_f.

The above can be extended to Λ: considering an element $x = x_0 + ux_1$ with $x_i \in \mathcal{O}_L$, $i = 0, 1$, we let the canonical embedding extend coefficient-wise for $\sigma_{\mathcal{A}}(x) := (\sigma_L(x_0), \sigma_L(x_1))$ and find that

$$V_\Lambda = \begin{pmatrix} \sigma_L(\overrightarrow{p}) & \mathbf{0} \\ \mathbf{0} & \sigma_L(\overrightarrow{p}) \end{pmatrix}$$

sends $\mathrm{coeff}(x) = (\mathrm{coeff}(x_0), \mathrm{coeff}(x_1))$ to $\sigma_{\mathcal{A}}(x)$. The singular values of this matrix are simply the singular values of $\sigma_L(\overrightarrow{p})$ multiplied by $\sqrt{2}$. As before, if $\ell = \mathrm{poly}(n)$, we find that the singular values of the above are polynomial in n, and similarly for V_Λ^{-1}.