

Imperial College of Science, Technology, and Medicine
Department of Electrical and Electronic Engineering

LWE Over Cyclic Algebras: A Novel Structure for Lattice Cryptography

Charles Everitt Grover

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Electrical and Electronic Engineering of Imperial College
London and the Diploma of Imperial College, November 2019

Abstract

The work done in this thesis is an introduction to the use of cyclic algebras in post quantum cryptography, with a particular focus on lattice cryptography and the Learning With Errors (LWE) problem. Algebraic variants of the LWE problem are some of the most promising candidates for quantum resistant public key and signature schemes. Since different variants come with their own strengths and weaknesses, adding a version over cyclic algebras provides another flavour to study.

The first part of the thesis establishes the hardness of the Learning With Errors Over Cyclic Algebras (CLWE) problem, basing its difficulty on short vector problems over certain structured lattices in the same manner as the analogous Ring and Module LWE problems. These structured lattices are applied in cryptography for the first time.

The second part of the thesis constructs concrete families of cyclic algebras suitable for use in practice. The requirements on dimensions of ambient spaces for lattice cryptography are considered and precise cyclic algebras are constructed using a collection of novel techniques.

The final part of this thesis explains carefully how to convert the theoretical construction of the CLWE problem into a form appropriate for real world cryptography. The problem is rephrased as a discrete problem, following which an example public key cryptosystem based on the CLWE problem is built in the style of other LWE based schemes.

Overall, the thesis introduces and thoroughly justifies the construction of the CLWE problem, the first instance of non-commutative ring cryptography based on lattices, and provides the basic requirements and functionalities needed to construct cryptographic primitives in these algebras. It also leaves behind a substantial quantity of new open questions regarding the use of cyclic algebras in more intricate cryptographic constructions.

Acknowledgements

I would like to take the time to express how thankful I am for those who have enabled me to complete this thesis. First of all, I find myself amazed by the constant enthusiasm that my supervisor, Cong Ling, has shown for our joint work, and I am pleased to have had the opportunity to work on his projects. I am also deeply appreciative of the academic insights provided by Roope Vekhalati and Antonio Campello during my forays into topics that I had not encountered before. To my family and close friends who have supported me over the last few years, thank you for your patience and tolerance. There have been too many members of CSP's 11:50am lunch group, of which I have enjoyed being a part, for me to thank individually, but I would like to make a few special mentions: Alex, for his sheer and outstanding omnipresence; Tricia, for her unique blend of unfaltering kindness and persevering realism; Wilhelm, for being particularly easy to share a desk with and moonlighting as my personal technical support team; Michael, because he asked; and Steph, for the understanding and encouragement provided to her fellow postgraduates. I am of the personal opinion that a good Acknowledgements section requires at least one pseudo-frivolous entry, and so I would like to offer my final thanks to the route setters at VauxEast and VauxWest bouldering gyms, whose excellent work has made my twice weekly climbing sessions a meditative experience that leave my mind refreshed for the rigors of research.

My research would not have been possible without the financial support of the ACE-CSR grant enabling me to undertake it, for which I am extremely grateful.

Statement of Originality

I hereby declare that all work included in this thesis is entirely my own, except where references are clearly stated.

Copyright Declaration

‘The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.’

“‘The time has come,’ the Walrus said,
‘To talk of many things:’
From: *The Walrus and the Carpenter*, by Lewis Carroll

Contents

Abstract	3
Acknowledgements	5
Declarations	7
1 Introduction	19
1.1 Motivation and Objectives	19
1.2 Contributions	28
1.3 Publications	30
2 Background Theory	31
2.1 Introduction	31
2.2 Lattices	32
2.2.1 Gaussian Distributions	33
2.2.2 Lattice Problems	34
2.3 The Learning With Errors Problem	35
2.3.1 Internal Reductions	37
2.4 Algebraic Number Theory	38

2.4.1	Basic Definitions	39
2.4.2	The Canonical Embedding	40
2.4.3	Relative Embeddings	41
2.4.4	Ideals and Ideal Lattices	42
2.5	Ring Learning With Errors	44
2.5.1	Primal Form, Normal Form, and RLWE as Structured LWE	45
2.6	Module LWE	47
2.6.1	MLWE as Structured LWE	48
2.7	Structured LWE in Cryptography	48
2.7.1	Comparing Ring and Module LWE	50
2.8	Cyclic Algebras	52
2.8.1	Matrix Representation	53
2.8.2	Orders and the Natural Order	53
2.8.3	Existence and non-Norm Elements	54
2.8.4	Order Ideals	55
2.8.5	Some Useful Ideals	57
2.8.6	Orders and Ideals as Integer Lattices	58
2.8.7	Gaussian Distributions Over Cyclic Algebras	58
3	Learning With Errors over Cyclic Algebras	60
3.1	Definition of CLWE	61
3.2	Lattice Problems Over Cyclic Algebras	63
3.2.1	Hardness of Order-Ideal Lattice Problems	64

3.3	Design Choices for CLWE	65
3.3.1	Relation to Module LWE	65
3.3.2	Natural Orders vs. Maximal	66
3.3.3	Security from Two Number Fields	68
3.3.4	Requiring the Division Property	68
3.3.5	CLWE as a One-Sided Problem	70
3.4	Technical Lemmas	71
3.5	Hardness of Search CLWE	73
3.5.1	BDD to Search CLWE	75
3.6	Search To Decision Reduction	78
3.6.1	Restricting the Secret Space	80
3.6.2	Outline of the Reduction	83
3.6.3	CLWE in $\mathcal{R}_i$	83
3.6.4	Hybrid CLWE and Search-Decision	85
3.6.5	Worst-Case to Average-Case Decision Reduction	87
3.6.6	Non-Applicability of the Direct to Decision Technique	88
4	Concrete Algebras for CLWE	93
4.1	Impossible Algebras	95
4.2	Suitable Algebras for Cryptography	96
4.2.1	Valid Extensions of Prime Power Cyclotomics	96
4.2.2	Constructions Using Compositums of Fields	98
4.3	Sample Parameters	99

4.3.1	Two-Power Cyclotomic K	100
4.3.2	Three-Power Cyclotomic K	101
4.3.3	Fields Using Compositum Techniques	101
4.4	Extensions Where q Splits Completely	102
4.4.1	Naive Method	103
4.4.2	Quadratic Case	104
4.4.3	Quartic Fields	105
5	CLWE in Cryptography	107
5.1	Discrete CLWE with Small Secrets	108
5.1.1	Primal Form	108
5.1.2	Normal Form	110
5.2	Example CLWE Cryptosystem	112
5.3	Operational Complexity in Cyclic Algebras	115
5.3.1	Algorithm for Multiplication in Cyclic Algebras	117
5.3.2	The Rings $\mathcal{R}_i$	118
5.3.3	Complexity of the CRT Style Isomorphism	120
5.3.4	Fast Cryptography When q Splits Completely in L	121
5.3.5	Generalizing to non-Split q and Arbitrary Bases	127
5.3.6	Generalizing to Other Centers	129
5.4	Attacking non-Division Algebras	130

6 Conclusion	135
6.1 Summary of Thesis Achievements	135
6.1.1 Security of CLWE	136
6.1.2 Constructing Appropriate Algebras	139
6.1.3 CLWE in Cryptography	140
6.2 Discussion and Future Work	141
6.2.1 Full Hardness of the Decision Problem	142
6.2.2 Fleshing Out CLWE	143
6.2.3 RLWE Attributes in CLWE	147
6.2.4 More Cryptographic Functionality	149
6.3 Concluding Remarks	150
Bibliography	150

List of Tables

4.1 Algebras Constructed in Section 4.3. 100

List of Figures

1.1	MLWE Matrix Over $\mathbb{Z}_q$	26
1.2	Security vs. Efficiency Trade-off for LWE Assumptions.	27
2.1	Matrix-Vector Representation of m LWE Samples.	36
2.2	Reductions for LWE. The Bold Arrow Denotes a Quantum Step.	37
2.3	Construction of a Cyclic Algebra.	52
4.1	Cyclic Extensions Between L and K	98
5.1	Depiction of Multiplication Algorithm.	118

Chapter 1

Introduction

Dark times lie ahead of us and there
will be a time when we must choose
between what is easy and what is right.

Albus Dumbledore, *Harry Potter and
The Goblet Of Fire* by J.K. Rowling

1.1 Motivation and Objectives

A permanent fixture of technological forecasting over the past twenty years has been the claim that the existence of reliable, scalable, quantum computers is a development which humanity must be ready to contend with in the near future. Although the details of this claim are often sparse, the relative success of modern quantum computing projects, such as those of IBM [Cas17] or D-Wave [Gib17], indicate that we, as engineers, must prepare ourselves for the reality of the quantum era sooner rather than later. Although understanding the exact capabilities of a truly quantum device is a science in its fledgling state, many fields within engineering and the wider scientific community can approach such machines with nothing but enthusiasm; they will most likely present some novel functionalities, such as in the field of optimization [NCS⁺17], that seem impossible on today's classical computing devices.

Unfortunately, cryptographers are forced to view such devices with a mix of justifiable enthusiasm and the discipline's trademark healthy scepticism. As is often the way with technological upgrades, the first thing learnt to do with them is to out-mode the ideas of their predecessors. In this case, the most widely celebrated quantum algorithm already known is that of Peter Shor [Sho94], an efficient algorithm for computing integer factorization. Practical implementations of Shor's algorithm would essentially invalidate RSA encryption [RSA78], based as it is on the difficulty of computing the factors of $n = pq$ for large hidden primes p and q . Furthermore, the Quantum Fourier Transform (QFT), perhaps the most remarkable building block of Shor's algorithm, can be applied to attacks on other common 'Pre-Quantum' cryptographic protocols such as Diffie-Hellman key exchange, since this relies on the hardness of the discrete logarithm problem which can also be solved via the QFT [Joz01].

Thus, the unprepared cryptographer would find themselves extraordinarily vulnerable to the first wave of functioning and scalable quantum devices, and a reliance on such foolish practitioners would spell disaster for the privacy of citizens, the proprietary dealings of corporations, and the secret workings of governments. Consequentially, much of the focus of modern cryptography is on what can be done to preemptively protect against quantum attacks. At first glance this may look an impossible task, since the hype surrounding quantum computers prescribes them with a somewhat mystical, enigmatic, nature that makes today's computers look helpless in response. However, as is the way of scientific research one can part the veil a little and begin to understand that no technology is all powerful. Whilst work on quantum algorithms such as that of Shor's provides insight into what a quantum device can do, along the way one can begin to form contrapositive ideas into what they will not be able to do as well. It is here that cryptographers who desire to count themselves as among the prepared may begin to formulate plans to prepare for the advent of quantum computing.

Of course, since the rumblings of a soon-to-exist quantum computer are no longer a new development, cryptographers have had time to consider how they might modernize in the wake of this new threat. Perhaps remarkably, the most common symmetric cryptographic scheme, AES [DR99], does not appear to need much in the way of updating in a quantum era bar perhaps a slight paradigm shift in what one considers a large enough security parameter. Certainly, one

should not expect to see a devastating attack or some other invalidation that would result in a need to find a replacement for information transfer between two parties with a shared key over an insecure channel.

However, the symmetric key cryptography performed by AES forms only half of the typical protocol for secure information transfer over the internet. After all, first the communicating parties must find a way to agree on said key, and if they were able to do this over a private channel then what use would they have for cryptography in the first place! Thus, attention turns to public key cryptography, the go-to method for exchanging symmetric keys in protocols such as TLS [Res15]. RSA encryption is the canonical example of public key techniques, and it is here the predicted devastating attacks by quantum devices present themselves in full force. Alternative techniques for Public Key Cryptography, such as those based on elliptic curves, also find themselves vulnerable to a quantum device [RNSL17]. Thus, a central problem of post-quantum cryptography and in part a motivating question for this thesis arises: what techniques can replace these public key techniques, and how can cryptographers be confident that the methods are secure against quantum computers?

In order to nudge the research community towards tackling this question, the immediacy of answering it has been made apparent by NIST, the National Institute of Standards and Technology of the United States of America. In 2016, they announced the commencement of a process [NIS16] to standardize the development and deployment of post-quantum cryptography, together with an open call for submissions from the wider cryptographic community. NIST ran similar processes in the past to standardize Symmetric Cryptography [NIS97] and Hash functions [NIS07], which resulted in the widely used AES and SHA-3 primitives respectively. Although NIST has been keen to emphasize that the “post-quantum standards development process should not be treated as a competition”, it is reasonable to expect that the ‘winner(s)’ of this process will be widely adopted in the manner of AES and SHA-3, and that this process represents initial movement towards solving the problem of quantum-resistant cryptography. As of Autumn 2019, the process is currently in its second round and there are 17 remaining public-key encryption¹ candidates as well as 9 remaining digital signature schemes [NIS19].

¹Technically, these schemes are Key Encapsulation Mechanisms (KEMs), but in this case the distinction is

Partitioning them by their underlying mathematical technique, 12 are schemes based on lattice cryptography, 7 on various forms of code based cryptography, 4 on multivariate methods, 2 on hash functions, and 1 on supersingular isogenies.

Almost half of the remaining candidates of the NIST process are lattice schemes, with particular over-representation amongst the public-key submissions, where 9 remain. Lattice based cryptography dates back to the work of Ajtai in the 1990's, who developed techniques for generating hard instances of famous lattice problems [Ajt96], such as the Shortest Vector Problem (SVP), and consequentially, along with Dwork, produced a public-key scheme with the difficulty of hard lattice problems as the underlying security guarantee [AD97]. Lattice problems such as SVP are particularly appealing for cryptographers since they are natural computational questions with deep implications in many disciplines of mathematics, and consequentially have been well studied for decades (see [LvdPdW12] for a comprehensive summary). To date, no efficient algorithms that solve these lattice problems with any great degree of efficiency exist, and the years of scrutiny they have been subject to lend some confidence to the notion that such problems will remain computationally intractable for the foreseeable future. Furthermore, despite substantial research efforts on attacking lattices with quantum algorithms (see e.g. [BS15], [LMvdP15], [BS16], [CDPR16]), appropriate lattice problems are yet to be substantially compromised in a quantum setting. While one must remain open to the fact that we have only scratched the surface of what quantum computers can do, for the time being it appears that basing post-quantum public-key cryptography on the hardness of SVP and related problems seems as salient a choice as any other.

One can divide the bulk of lattice based cryptography into two categories. The first category dates back to 1996 and is comprised of schemes based on the mysteriously named NTRU system, first published a couple of years later in [HPS98]. Formally, NTRU is based on computation over polynomial rings, but these rings are intimately related with lattices and much of the cryptographic and cryptanalytic work done on NTRU uses the lattice relationship. Although NTRU does not permit a reduction to any standard lattice problems such as SVP, it has robustly stood the test of time and rebuffed the bulk of cryptanalysis thrown at it, at least as largely semantic.

long as one displays a degree of care when choosing parameters (see e.g. [CS97], [ABD16], and [KF17]). Three schemes based on NTRU survived the first round of the NIST process, two of which are public key schemes. One of these is essentially the classical NTRU scheme and thus is just called NTRU [CDH⁺18]. The other is known as NTRUPrime [BCLvV17], and contains some relatively substantial updates based on modern developments in lattice cryptography.

On the other hand, there is cryptography based on the Learning with Errors (LWE) problem, introduced by Regev in 2005 [Reg05]. In Regev’s seminal work he demonstrated that the problem of solving the approximate matrix-vector equation $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ over a finite field $\mathbb{Z}_q$ for public matrix $\mathbf{A}$, unknown secret $\mathbf{s}$, and small error $\mathbf{e}$ is at least as hard as solving SVP in the worst-case for reasonable approximation factors. To add further credence to the idea that solving the LWE problem is hard, Regev observed that it corresponds to the coding theoretic problem of decoding random linear codes, for which no fast general algorithms are known despite it being an object of substantial study. There are 10 NIST second round schemes based on the hardness of variants of the LWE problem, although only one, FrodoKEM [ABD⁺18], is based on the hardness of the initial LWE problem itself.

So what exactly are the remaining 9 second round lattice schemes based on? To understand this one must consider the limitations of the LWE problem. A scheme based on LWE might store the matrix $\mathbf{A}$ in its public key, a large and uniformly random matrix over a finite field². Storing such a matrix does not come at a cheap cost; hence, a natural question to consider is how one might squash the cost of storing the matrix $\mathbf{A}$. The obvious method would be to replace the uniformly random structure of $\mathbf{A}$ with a structure that is somehow predictable or repetitive. However, one can not just place some structure on $\mathbf{A}$ and assume the scheme is secure; for example, in the trivial setting where $\mathbf{A}$ is a random scaling of the identity matrix one can imagine simple algorithms to recover $\mathbf{s}$. Instead, one needs to structure the random matrix $\mathbf{A}$ in such a way as to maintain a satisfactory provable security guarantee whilst also streamlining storage requirements.

The first steps in this direction were taken by Lyubashevsky, Peikert, and Regev in 2010

²In fact, the NIST submissions using LWE typically squash the storage of $\mathbf{A}$ by just storing a seed for a PRG generating it, but the intuition here is still useful.

[LPR10], although a concurrent work of Stehlé, Steinfeld, Tanaka, and Xagawa contains a similar line of reasoning [SSTX09]. In [LPR10] the authors introduced the Ring Learning with Errors (RLWE) problem, a prominent variant of LWE whereby the noisy equation to solve is $b = a \cdot s + e$ for unknown s , short e , and uniformly random a all taken from certain subspaces of a number field K . This equation is more commonly expressed in its ‘primal’ form, where $a, s, e \in R_q := \frac{\mathbb{Z}_q[x]}{x^N+1}$ for a prime modulus q and N a power-of-two, and it is from here we derive a structured instance of the classical LWE problem. Multiplication in the ring R_q can be described by a matrix-vector operation as follows: let $\mathbf{s}$ and $\mathbf{e}$ denote the coefficient vectors of s and e respectively and let a_i denote the i^{th} coefficient of a . Then, the coefficient vector $\mathbf{b}$ of b can be computed as

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

where the matrix on the left, typically denoted $\text{rot}(a)$ and known as a negacyclic matrix, has taken the place of the uniformly random matrix $\mathbf{A}$ from the LWE problem. Recalling that $\mathbf{A}$ is likely to represent the public key in an encryption scheme, it follows that storing the now structured public key $\mathbf{A} = \text{rot}(a)$ requires storing only the N elements of the first column rather than all N^2 entries of the matrix. The NewHope scheme, currently in the second round of the NIST process [ADPS16], is based on the primal variant of the RLWE problem, sometimes referred to as the Polynomial Learning With Errors (PLWE) problem. Additionally, this scheme was briefly deployed on the internet by Google as a proof of concept implementation for post-quantum public key encryption [Bra16], and it appeared to be essentially functional.

By itself, merely presenting a way to structure the matrix $\mathbf{A}$ is not enough to become a cryptographic protocol. In [LPR10] the authors also provide a security reduction from the RLWE problem to short vector problems over so-called ideal lattices. Ideal lattices are structured lattices derived from the canonical embedding of ideals in the ring of integers $\mathcal{O}_K$ of a number field K . Intuitively, one can view these lattices as being generated by vectors with a cyclic structure

corresponding to the shape of $\text{rot}(a)$, although this hides some technical details. The structure imbued by the canonical embedding results in a lattice problem that is at most as hard as the generic SVP problem, but despite great efforts to exploit this structure there does not appear to be a substantial gap between the hardness of ideal lattice problems and their unstructured equivalents. Indeed, the current technique for attacking instances of the RLWE problem is to treat them as unstructured instances of the LWE problem and attack that, as in the CoreSVP method of [ADPS16]. However, it is still prudent to be cautious, and following algorithms utilizing the algebraic structure in attacks on related ideal lattice problems [CDW17] it seems plausible that in the coming years techniques may develop that widen the security gap between RLWE and classical LWE.

Although RLWE represents an elegant solution to the problem of storage space, it too is not without issues. To provide both security and efficiency the dimension N is typically restricted to being a power-of-two, hampering the flexibility of ring based schemes. NewHope, for example, targets varying levels of security with $N = 512$ and $N = 1024$ but cannot achieve any intermediate dimension while maintaining that N is a power-of-two. Since higher dimension typically means slower functionality, this inflexibility can cause schemes to become a little inefficient. Consequentially, Langlois and Stehlé introduced Module LWE (MLWE) in 2015 [LS15], in part to allow for targeting of a wider range of dimensions. Here the ring of dimension N over $\mathbb{Z}_q$ is replaced by a module M of rank d over the ring R_q in dimension n and the noisy equation to solve is $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ for public $\mathbf{A} \in M^d$, with the secret $\mathbf{s}$ and short error $\mathbf{e}$ both contained in M . Denoting the total dimension by $N = nd$ it follows that the MLWE problem strictly generalizes the classical and ring variants of the LWE problem: $n = 1, d = N$ is precisely the classical LWE problem whereas $n = N, d = 1$ is the ring variant. Once more, for both security and efficiency the ring dimension n is set to be a power-of-two, but in this case the variable rank d allows for more flexibility in terms of total dimension N . For example, $n = 256$ and $d = 3$ achieves total dimension 768, lying between the two dimensions of NewHope. Again, there is a second round NIST process candidate based on this security problem: CRYSTALS-Kyber [BDK⁺18] is an MLWE based scheme with fixed ring dimension $n = 256$ and flexible module rank d to target differing levels of security.

Returning to the representation of all these problems as matrix-vector equations over a finite field $\mathbb{Z}_q$, we find that the MLWE problem presents itself in a slightly unusual fashion. With d samples, the module equation, where entries are in R_q , can be written as

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

for $a_{i,j} \in R_q$, each of which is a polynomial corresponding to a $\mathbb{Z}_q$ -vector of dimension n . Here, we may apply the $\text{rot}(a)$ operation coordinatewise to each entry of $\mathbf{A}$ and correspondingly express each element of $\mathbf{s}$, $\mathbf{e}$, and $\mathbf{b}$ in their coefficient vectors to recover the $\mathbb{Z}_q$ matrix-vector equation $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$, where the matrix $\mathbf{A}$ takes the following form:




Figure 1.1: MLWE Matrix Over $\mathbb{Z}_q$.

a collection of d^2 blocks, each of which is a negacyclic matrix in dimension n . In terms of storing such a public matrix $\mathbf{A}$, each negacyclic block can be stored via its first column but each block must be stored separately. The overall storage cost is therefore $nd^2 = Nd$ elements of $\mathbb{Z}_q$, an improvement on an LWE public key but a factor of d worse than a corresponding RLWE key. Here the trade-off between efficiency and security typical of cryptography appears,

since although the efficiency of an MLWE scheme essentially decreases as the module rank d increases relative to dimension N , security effectively increases. The paper [LS15] introducing MLWE provides a security reduction from the MLWE problem to the Module Lattice-SIVP (Shortest Independent Vectors Problem), which is no easier than the corresponding problem on ideal lattices, so it seems that MLWE has advantages over RLWE on two fronts: dimensional flexibility and improved security, which come at a cost in efficiency of operations and storage requirements.

The remaining second round candidates for the NIST post-quantum process indicate that there are viable lattice techniques other than these three LWE assumptions: ThreeBears [Ham17] bases itself on a novel ‘integer module’ variant of LWE, NTRU and NTRUPrime use the techniques of NTRU, and Round5 [BBF⁺18] and Saber [DKRV18] use rounded variants called (ring) Learning with Rounding that replace the additive error term $\mathbf{e}$ with ‘error’ induced by a deterministic rounding step modulo some other prime p . However, understanding the concrete security and efficiency gaps between the classical, module, and ring variants of the LWE problem remains a central area of research in lattice cryptography, although the works of [RSW18] and [PP19] indicate that the security of most LWE variants are asymptotically equivalent in polynomial time.




Figure 1.2: Security vs. Efficiency Trade-off for LWE Assumptions.

The heavy algebraic structure in RLWE is an area of security concern, and the development of quantum algorithms utilizing the field structure in the last few years indicates that one must be careful with non-standard security problems (see for example [CGS14], [BS16], [CDPR16]). Additionally, there is some concern that the power-of-two dimension rings, popular for their

efficient multiplication algorithms and associated security reductions, may be more vulnerable than first thought due to the remarkable structure of the Galois group of the underlying number fields, known as cyclotomic fields. In fact, this worry motivates certain choices in the NTRUPrime NIST candidate (see Appendix A of [BCLvV17]). Many of the associated quantum algorithms rely, either explicitly or implicitly, on the commutative nature of the algebraic number field K and the well-known structure of its unit group, such as in [EHKS14] or [BS16]. Conversely, this structure is not so clearly present in the module case, from which a plausible assertion can be made that schemes based on RLWE could one day end up compromised without substantial concrete effect on corresponding MLWE schemes, even though in polynomial time their security turns out to be equivalent (see [AD17]). On the other hand, RLWE schemes are typically more efficient than the corresponding MLWE schemes, an advantage one would not want to lose if the RLWE schemes turn out to be sufficiently secure. This motivates the central research question of this thesis: does there exist an LWE style assumption that combines some of the benefits of RLWE and MLWE?

1.2 Contributions

Although MLWE generalizes RLWE in the sense that at $d = 1$ the MLWE problem is precisely the RLWE problem, the jump from module rank 1 to 2 represents a large theoretical leap: the deep algebraic structure of the ring is pushed out by the non-trivial module rank. Furthermore, as the module rank increases the cyclical storage structure is broken down into smaller and smaller blocks, each of which must be stored separately as a ring element. In this thesis, we present a novel LWE problem, called the Cyclic Learning With Errors (CLWE) problem, in an attempt to reproduce the module over a number field structure of MLWE with an algebra over a number field structure, in order to maintain the efficiency gains derived from algebraic structure whilst mitigating some of the potential security weaknesses of the ring based schemes.

First, we construct the CLWE assumption using cyclic algebras over number fields, a tool of increasing popularity in the field of coding theory (see e.g. [ORBv06], [OBV07], [VHLR09],

[LVL18] among others). Cyclic algebras are remarkable for permitting a matrix representation of left multiplication: for a well chosen extension of number fields L/K of degree d , multiplication in a cyclic algebra $a \cdot s$ can be written in the matrix-vector form $\phi(a) \cdot \mathbf{s}$, with

$$\phi(a) = \begin{pmatrix} a_0 & \gamma\theta(a_{d-1}) & \gamma\theta^2(a_{d-2}) & \dots & \gamma\theta^{d-1}(a_1) \\ a_1 & \theta(a_0) & \gamma\theta^2(a_{d-1}) & \dots & \gamma\theta^{d-1}(a_2) \\ a_2 & \theta(a_1) & \theta^2(a_0) & \dots & \gamma\theta^{d-1}(a_3) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{d-1} & \theta(a_{d-2}) & \theta^2(a_{d-3}) & \dots & \theta^{d-1}(a_0) \end{pmatrix}$$

where each $a_i \in L$, γ is a fixed element of K , θ is a suitable K -automorphism of L , and $\mathbf{s}$ is the coefficient vector of s as a module element over L . Restricting to an appropriate discrete structure Λ_q known as the natural order, corresponding to the rings R_q used in ring and module LWE, we define a CLWE instance $b = a \cdot s + e$ for uniformly random public a , secret s , and small e all in Λ_q . Then, we recover a matrix-vector operation over appropriately chosen R_q of dimension nd by rewriting the equation as $\mathbf{b} = \phi(a) \cdot \mathbf{s} + \mathbf{e}$. Finally, by applying $\text{rot}(a)$ coordinatewise we recover a matrix-vector operation over $\mathbb{Z}_q$ of dimension $N = nd^2$ containing d^2 blocks of dimension nd . Since the first column of $\phi(a)$ uniquely determines the whole matrix, one need only store the first d blocks of $\phi(a)$. Since each $\text{rot}(a_i)$ operation streamlines storage by a factor of nd , we see that the whole $\mathbb{Z}_q$ -matrix can be stored in space $O(N)$, just as in RLWE, by storing its first column.

We provide a reduction from the CLWE problem to structured lattice problems: namely, the SVP over the structured lattices generated by embeddings of ideals of the natural order Λ_q , which we call order ideal SVP, corresponding to the ideal SVP problems underlying RLWE. We also demonstrate that an important security property common to LWE instances which was not discussed in this introduction also holds in the case of CLWE, namely a search-decision equivalence, although unfortunately in our case this comes at a loss in the size of the secret space for the search problem. We provide internal reductions on the decision problem, demonstrating worst-case to average-case security and primal and normal form reductions. We remark that CLWE is a true generalization of RLWE, in the sense that at $d = 1$ the CLWE problem is

precisely the RLWE problem in dimension n .

We further discuss and present substantial design choices for CLWE, with a view towards demonstrating what parameter sets may be suitable for practical applications. For example, we discuss a variety of constructions of cyclic algebras that achieve appropriate dimensions for cryptography, since the existence of cyclic algebras in a given dimension is a more complicated problem than that of number fields or modules. We explain how the choice of modulus q effects algorithms for computing the matrix-vector operations and present asymptotically efficient algorithms for multiplication in a cyclic algebra. We will also present alternative construction choices which we did not make in this thesis, discussing their shortcomings and what techniques one would need to instantiate CLWE with those design specifications. Finally, we present a sample cryptosystem whose security is based on the hardness of the CLWE problem.

1.3 Publications

Charles Grover, Cong Ling, and Roope Vehkalahti. Non-Commutative Ring Learning With Errors From Cyclic Algebras. Cryptology ePrint Archive, Report 2019/680, 2019.

Chapter 2

Background Theory

A straight line may be the shortest distance between two points, but it is by no means the most interesting.

The Doctor, *Doctor Who: The Time*

Warrior by Robert Holmes

2.1 Introduction

In this chapter the collection of fundamentals required to understand this thesis will be introduced. We will proceed in a broadly chronological manner in terms of appearance of concepts in cryptographic literature, beginning with a standard introduction to lattices, Gaussian distributions, and computational lattice problems that together underpin all modern lattice based cryptography. Following these purely linear algebraic concepts we introduce the definition of the classical LWE problem. We then discuss the necessary concepts from algebraic number theory required to understand Ring and Module LWE, including how one constructs lattices from algebraic structures, then rigorously define both assumptions. Finally, we move on to the mathematical notions that are previously unused in lattice based cryptography: we define cyclic algebras and present some well known theorems on their structure which we require for

our novel LWE construction.

2.2 Lattices

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

Definition 2.1. Given a lattice $\mathcal{L}$ in a space V endowed with a metric $\|\cdot\|$, the minimum distance of $\mathcal{L}$ is defined as $\lambda_1(\mathcal{L}) = \min_{\mathbf{v} \in \mathcal{L} \setminus \{0\}} \|\mathbf{v}\|$. Similarly, the n^{th} minima $\lambda_n(\mathcal{L})$ is the minimum length of a set of n linearly independent lattice vectors, where the length of a set of vectors $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ is defined as $\max_i(\|\mathbf{x}_i\|)$.

Definition 2.2. Given a lattice $\mathcal{L}$ in V , where V is endowed with an inner product $\langle \cdot, \cdot \rangle$, the dual lattice $\mathcal{L}^*$ is defined $\mathcal{L}^* = \{\mathbf{v} \in V : \langle \mathcal{L}, \mathbf{v} \rangle \subset \mathbb{Z}\}$.

Definition 2.3. Given a lattice $\mathcal{L}$ in V , a fundamental domain is any set $\mathcal{F}(\mathcal{L}) \subset V$ which contains precisely one representative of the coset $\mathbf{x} + \mathcal{L}$ for each $\mathbf{x} \in V$. That is, $|(\mathbf{x} + \mathcal{L}) \cap \mathcal{F}(\mathcal{L})| = 1$ for each $\mathbf{x} \in V$.

The relevance of the first definition is perhaps self evident, since in the introduction we remarked that the problem of finding short vectors is of central importance to lattice cryptography. The dual lattice is a central object of study in lattice theory, and will arise frequently in the formalization of LWE problems. The fundamental domain relates the discrete object that is a lattice to more continuous geometric properties, and can be best understood by the following example: for the integer lattice $\mathbb{Z}^n \subset \mathbb{R}^n$ the fundamental domain can be chosen as a hypercube, since for example $\mathcal{F}(\mathbb{Z}^n) := [-\frac{1}{2}, \frac{1}{2})^n$ can easily be seen to contain exactly one representative of $\mathbf{x} + \mathbb{Z}^n$ for any $\mathbf{x} \in \mathbb{R}^n$. This is also an example of a special case of a fundamental domain

known as the fundamental parallelepiped, denoted $\mathcal{P}(\mathcal{L})$, which allows us to define the notion of the volume of a lattice.

Definition 2.4. Given a lattice $\mathcal{L} \subset \mathbb{R}^n$ with basis $\mathbf{B}$, the (origin centered) fundamental parallelepiped $\mathcal{P}(\mathcal{L})$ is the continuous region $\mathbf{B} \cdot [-\frac{1}{2}, \frac{1}{2}]^n$. The volume $\text{Vol}(\mathcal{L})$ of $\mathcal{L}$ is the volume of $\mathcal{P}(\mathcal{L})$ as a subspace of $\mathbb{R}^n$.

It is well known that the volume of a lattice is invariant under the choice of basis $\mathbf{B}$, which follows immediately from the fact that if $\mathbf{B}_1$ and $\mathbf{B}_2$ are bases of the same lattice then there is a unimodular matrix $\mathbf{U}$ satisfying $\mathbf{B}_1 = \mathbf{B}_2 \mathbf{U}$.

2.2.1 Gaussian Distributions

Definition 2.5. For a vector space V with norm $\|\cdot\|$ and an $r > 0$ we define the Gaussian function $\rho_r : V \rightarrow (0, 1]$ by $\rho_r(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/r^2)$.

We 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 $\mathbf{b}_1, \dots, \mathbf{b}_n$ of V , for $\mathbf{r} = (r_1, \dots, r_n)$ a vector of positive reals, by sampling $x_1, \dots, x_n$ independently from the one dimensional Gaussian distributions D_{r_i} and outputting $\sum_{i=1}^n x_i \mathbf{b}_i$.

When sampling over a lattice $\mathcal{L}$ we require a discretized form of the Gaussian distribution. We define the discrete Gaussian distribution $D_{\mathcal{L},r}$ over $\mathcal{L}$ as outputting $\mathbf{x}$ with probability $\frac{\rho_r(\mathbf{x})}{\rho_r(\mathcal{L})}$ for each $\mathbf{x} \in \mathcal{L}$, where $\rho_r(\mathcal{L}) := \sum_{\mathbf{x} \in \mathcal{L}} \rho_r(\mathbf{x})$. Note that our defined Gaussian distributions are centered at 0, but one may replace these with more general distributions $D_{\mathbf{c},r}$ centered at arbitrary $\mathbf{c} \in V$ by defining $\rho_{\mathbf{c},r}(\mathbf{x}) = \rho_r(\mathbf{x} - \mathbf{c})$.

An important lattice quantity known as the smoothing parameter was introduced in [MR07]. The motivation for the name is provided by Lemma 2.7 following the definition.

Definition 2.6. For a lattice $\mathcal{L}$ and $\varepsilon > 0$ the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ is defined as the smallest $r > 0$ satisfying $\rho_{1/r}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \varepsilon$.

The following is a special case of [MR07], Lemma 4.1.

Lemma 2.7. *For a lattice $\mathcal{L}$ over $\mathbb{R}^n$, $\varepsilon > 0$, $r \geq \eta_\varepsilon(\mathcal{L})$, and $\mathbf{x} \in \mathbb{R}^n$, the statistical distance between $(D_r + \mathbf{x}) \bmod \mathcal{L}$ and the uniform distribution modulo $\mathcal{L}$ is bounded above by $\varepsilon/2$. Equivalently, $\rho_r(\mathcal{L} + \mathbf{x}) \in [\frac{1-\varepsilon}{1+\varepsilon}, 1] \cdot \rho_r(\mathcal{L})$.*

Informally, the smoothing parameter can be thought of as follows: for r exceeding the smoothing parameter, the Gaussian distribution modulo $\mathcal{L}$ is statistically close to uniform on the fundamental domain of $\mathcal{L}$.

2.2.2 Lattice Problems

Computational problems on lattices represent the foundations of the security of existing LWE assumptions and will do so for our novel CLWE assumption as well. The archetypal lattice problem is the shortest vector problem, that of finding a lattice vector $\mathbf{x}$ achieving the first minimum $\lambda_1(\mathcal{L})$ of the lattice. Similarly, the shortest independent vectors problem is that of finding a set of n linearly independent vectors whose maximum length is $\lambda_n(\mathcal{L})$, the n^{th} minima of $\mathcal{L}$. In practice, we will be more interested in approximate variants of these problems.

Definition 2.8. Let $\|\cdot\|$ be some norm on $\mathbb{R}^n$ and let $\xi := \xi(n) \geq 1$. Then, the approximate Shortest Vector Problem (SVP_ξ) on input a lattice $\mathcal{L}$ is to find some non-zero vector $\mathbf{x}$ such that $\|\mathbf{x}\| \leq \xi \cdot \lambda_1(\mathcal{L})$.

Definition 2.9. Let $\|\cdot\|$ be some norm on $\mathbb{R}^n$ and let $\xi := \xi(n) \geq 1$. Then, the (approximate) Shortest Independent Vectors Problem (SIVP_ξ) on input a lattice $\mathcal{L}$ is to find n linearly independent non-zero vectors $\mathbf{x}_1, \dots, \mathbf{x}_n$ such that $\max_i(\|\mathbf{x}_i\|) \leq \xi \cdot \lambda_n(\mathcal{L})$.

The computational hardness of these problems forms the backbone of lattice cryptography, since to date there are no algorithms that efficiently solve these problems for polynomial approximation factors ξ . For our needs it is sufficient to know that these are ‘hard’ to solve computationally, and the precise qualities of state of the art algorithms for the problems will not play a part in this thesis. Another natural lattice problem is that of lattice decoding, known

as the Closest Vector Problem: given a point $\mathbf{c}$ in ambient space $\mathbb{R}^n$, find the closest $\mathbf{x} \in \mathcal{L}$ e.g. the lattice point $\mathbf{x}$ such that $\|\mathbf{x} - \mathbf{c}\|$ is minimized. In LWE security proofs, this problem arises with the additional guarantee that the vector $\mathbf{c}$ is not too far from the lattice, typically known as bounded distance decoding.

Definition 2.10. Let $\|\cdot\|$ be some norm on $\mathbb{R}^n$, let $\mathcal{L}$ be a lattice, and let $\delta < \lambda_1(\mathcal{L})/2$. Then, the Bounded Distance Decoding problem ($\text{BDD}_{\mathcal{L},\delta}$) on input $\mathbf{y} = \mathbf{x} + \mathbf{e}$ for $\mathbf{x} \in \mathcal{L}$ and $\|\mathbf{e}\| \leq \delta$ is to compute $\mathbf{x}$, or equivalently $\mathbf{e}$.

Though there are many other well studied computational problems on lattices, these three are the ones that arise most frequently in the context of lattice cryptography, particularly in the context of security reductions.

2.3 The Learning With Errors Problem

We are now ready to define the classical LWE problem, introduced by Regev in [Reg05]. LWE problems come in two varieties, search and decision, both derived from an LWE distribution. This can be defined over a continuous or discrete space, but we will present only the discrete form in this section due to its conceptual simplicity.

Definition 2.11. For a secret vector $\mathbf{s} \in \mathbb{Z}_q^n$ and an error distribution χ over $\mathbb{Z}_q$ the LWE distribution $A_{\mathbf{s},\chi}$ is sampled from as follows: sample $\mathbf{a} \leftarrow \mathbb{Z}_q^n$ uniformly at random, sample $e \leftarrow \chi$, and output $(\mathbf{a}, b) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e \bmod q) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$.

Typically χ will be a discrete, 0-centered, Gaussian distribution with low variance. The foundations of this distribution as inducing a hard problem can be seen by its relation to the famous Learning Parity with Noise (LPN) problem [BJMM12], since if $q = 2$ and χ is Bernoulli this distribution is exactly the one used in LPN. We define two related computational problems.

Definition 2.12. The search LWE problem, denoted $\text{LWE}_{\mathbf{s},\chi}$, is to find $\mathbf{s}$ given access to a collection of independent samples $(\mathbf{a}_i, b_i)$ drawn from $A_{\mathbf{s},\chi}$ for unknown, uniformly random, secret $\mathbf{s}$.

Definition 2.13. The decision LWE problem, $\text{D-LWE}_{\mathbf{s},\chi}$, is as follows: upon being given access to samples $(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ where either each sample was taken from the distribution $A_{\mathbf{s},\chi}$ for random $\mathbf{s}$ or each taken from the uniform distribution over $\mathbb{Z}_q^{n+1}$ to decide which is the case with non-negligible advantage

Formally, both problems should be parameterized by the number m of available samples, but the choice of m is typically dependent upon the application so we drop it for representational simplicity. Often in practical applications the number of samples is set so that $m = n$, and the n samples $(\mathbf{a}_i, b_i)$ are concatenated together into the matrix-vector equation $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ discussed in the introduction.




Figure 2.1: Matrix-Vector Representation of m LWE Samples.

For applications in cryptography, the hardness of the decision problem is preferred to that of the search problem. Assuming that the decision problem is hard implies that LWE samples are computationally indistinguishable from uniform, so intuitively an LWE sample can be used to hide a message m as an element of $\mathbb{Z}_q^n$ by adding it to b ; then since the LWE samples are assumed to be indistinguishable from uniform an eavesdropper who does not know $\mathbf{s}$ cannot discern any information about m from $(\mathbf{a}, b + m)$. In [Reg05], Regev demonstrated a reduction from the approximate SVP to the search LWE problem, implying that $\text{LWE}_{\mathbf{s},\chi}$ is at least as hard as approximate SVP. This proof follows a sequence of shorter reductions: first, the approximate SVP is reduced to the problem of sampling a discrete Gaussian of narrow variance over a lattice $\mathcal{L}$, where intuitively sampling from a sufficiently narrow Gaussian should output a vector whose

norm is reasonably short compared to the first minima. Then, a quantum algorithm reduces the problem of sampling from a narrow Gaussian to that of solving the BDD problem on the lattice. This is an iterative step which uses a BDD solver to output samples from progressively narrower Gaussian distributions. Finally, a transformation maps an instance of the BDD problem to an appropriate instance of the LWE problem, reducing the BDD problem to that of LWE.




Figure 2.2: Reductions for LWE. The Bold Arrow Denotes a Quantum Step.

To visualize the transformation between the LWE and the BDD problems, consider the q -ary lattice $\mathcal{L}(\mathbf{A})$ generated by the rows of $\mathbf{A}$, $\mathcal{L} = \{\mathbf{A}^t \mathbf{x} \mid \mathbf{x} \in \mathbb{Z}_q^n\} + q\mathbb{Z}^n$. Here $\mathbf{b}$ is close to one element of $\mathcal{L}(\mathbf{A})$, corresponding to the secret vector $\mathbf{s}$, and one can see the intimate relationship between lattice problems and the LWE problem by viewing it as a decoding problem on this lattice. In particular, since we can bound the size of the Gaussian error terms with high probability it is clear that this is almost always a bounded distance decoding problem.

2.3.1 Internal Reductions

So far, this sequence of reductions only establishes the hardness of the search variant of the LWE problem. In order to demonstrate the hardness of the decision problem [Reg05] provides a reduction from the search problem to the decision problem. This reduction proceeds by effectively ‘guessing’ each coordinate of the secret $\mathbf{s}$, verifying the guesses independently using

an oracle for the decision problem until each coordinate is found. Since the decision problem is trivially at most as hard as the search problem, this establishes the polynomial time equivalence of the two variants. We will see later that the methodology of [Reg05] can be used to show the hardness of the search version of more structured LWE problems, whereas the search-decision equivalence becomes more complicated. Finally, [Reg05] also provides a worst-case to average-case reduction for the LWE problem. Formally, this states that being able to solve the D-LWE problem for a non-negligible proportion of possible secrets $\mathbf{s}$ implies one can solve it for any secret $\mathbf{s}$, but in practice it means that a random instance of the decision LWE problem can be considered hard as long as some instances are hard, a necessity in the case where one wishes to implement cryptography with randomly chosen $\mathbf{s}$.

Another important internal security reduction comes from [ACPS09], in which the authors show that the hardness of the LWE problem $\text{LWE}_{\mathbf{s},\chi}$, where $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ is uniformly random, implies the hardness of the problem $\text{LWE}_{\mathbf{s}',\chi}$ when each coordinate of $\mathbf{s}'$ is sampled independently from the error distribution χ . When the secret is chosen by taking each coordinate independently from the error distribution we say that the LWE problem is in normal form, and the normal form of the LWE problem is typically used in practical settings since having a small $\mathbf{s}$ results in more efficient computations and a particularly compact cryptosystem.

2.4 Algebraic Number Theory

We were able to discuss the LWE problem in the previous section using only fundamental concepts from linear algebra. Before we can understand its structured variants we must first present a thorough grounding in the basics of algebraic number theory. Here we also introduce the theory required to understand the construction of cyclic algebras, which we will come to later in this chapter.

2.4.1 Basic Definitions

Definition 2.14. 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 dimension as a field extension.

When defining 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 2.15. Let L/K be an extension of number fields of degree n . The Galois group of L over K is the group $\text{Aut}(L/K)$ of automorphisms of L that fix K . We say that the extension is Galois if the subfield of L fixed by $\text{Aut}(L/K)$ is exactly K .

Note that this definition does not preclude the case where K is the trivial extension $\mathbb{Q}$, in which the Galois group of L is precisely the group of automorphisms of L that fix the rationals. We define a cyclic Galois extension L/K to be a Galois extension such that the Galois group of L over K is the cyclic group generated by some element θ of degree $d := [L : K]$. A central object in the algebraically structured variants of the LWE problem is the ring of integers of a number field, which will take the place of the ambient space $\mathbb{Z}_q$ used in LWE. In the following definition we use the usual definition of the minimal polynomial of an element of $\alpha \in \mathbb{C}$ over $\mathbb{Q}$; the unique irreducible monic polynomial $f(x) \in \mathbb{Q}[x]$ of least degree satisfying $f(\alpha) = 0$.

Definition 2.16. Given a number field K its ring of integers $\mathcal{O}_K$ is the ring consisting of those elements of K whose minimal polynomials over $\mathbb{Q}$ lie in $\mathbb{Z}[x]$.

In cryptographic literature it is common to denote $\mathcal{O}_K$ simply by R . However, since we will often be working with a pair of number fields we stick to the notation $\mathcal{O}_K$ and $\mathcal{O}_L$ to avoid confusion. It is easy to check that if L/K is an extension of number fields then $\mathcal{O}_L \cap K = \mathcal{O}_K$. We will refer to elements of the ring of integers of a number field as integral.

2.4.2 The Canonical Embedding

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

Definition 2.17. For $\beta \in K$ we define the trace and norm of β as $\text{Tr}_{K/\mathbb{Q}}(\beta) = \sum_{i=1}^n \sigma_i(\beta)$ and $\text{Norm}_{K/\mathbb{Q}}(\beta) = \prod_{i=1}^n \sigma_i(\beta)$ respectively. We will sometimes drop the subscript and just write e.g. $\text{Tr}(\beta)$ when the number field is clear from context.

Definition 2.18. Let $K = \mathbb{Q}(\alpha)$ be a number field of degree $n = r_1 + 2r_2$. The canonical embedding σ is the ring homomorphism $\sigma : K \rightarrow \mathbb{R}^{r_1} \times \mathbb{C}^{2r_2}$ defined by

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

Formally, σ maps into the space

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

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 $r_1 < j \leq r_1 + r_2$, then use the well-defined ℓ_p norm induced by viewing H as a subset of $\mathbb{C}^n$. Observe that, since each σ_i is a homomorphism, 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 $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R} \cong H$

rather than over K itself. 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 when necessary to avoid confusion.

2.4.3 Relative Embeddings

In the case of an extension L of a number field K it is sometimes more convenient to apply a different order to its embeddings induced by extending embeddings of K to those of L . Given a tower $L/K/\mathbb{Q}$ where $[K : \mathbb{Q}] = n$ and $[L : K] = d$ there are precisely n embeddings $\sigma_1, \dots, \sigma_n$ of K into $\mathbb{C}$. As long as $L/\mathbb{Q}$ is a Galois extension each of these can be extended to an embedding $\alpha_i : L \rightarrow L$ such that $\alpha_i|_K = \sigma_i$. However, these extensions are not unique, and it is easy to see that in the Galois case there are precisely $[L : K] = d$ choices for each α_i . In particular, in the case where L/K is a cyclic extension with Galois group generated by θ it holds that the composite automorphisms $\theta^j \circ \alpha_i(\cdot)$, $1 \leq j \leq d$, run through the d choices of α_i . Hence for a fixed choice of $\alpha_1, \dots, \alpha_n$ the nd automorphisms of L can each be uniquely represented by some $\theta^j \circ \alpha_i(\cdot)$. 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. We will make use of the ordering running through j later, such that the embeddings come in n blocks of length d with each block corresponding to the d L -extensions of one fixed K -embedding. We can also use the embeddings of L fixing K to define relative versions of the trace and norm.

Definition 2.19. Let L/K be an extension of number fields and let $\theta_1, \dots, \theta_d$ be the embeddings of L that fix K . Define the relative trace and norm as maps from L into K as $\text{Tr}_{L/K}(\ell) = \sum_i \theta_i(\ell)$ and $\text{Norm}_{L/K}(\ell) = \prod_i \theta_i(\ell)$ respectively.

Remark. Both of these maps are transitive in the sense that for a tower of finite extensions $M/L/K$ and $m \in M$ we have $\text{Tr}_{M/K}(m) = \text{Tr}_{L/K}(\text{Tr}_{M/L}(m))$ and $\text{Norm}_{M/K}(m) = \text{Norm}_{L/K}(\text{Norm}_{M/L}(m))$.

2.4.4 Ideals and Ideal Lattices

In order to construct lattices from number fields one applies the canonical embedding to an ideal $\mathcal{I}$ of $\mathcal{O}_K$, where as usual an ideal in $\mathcal{O}_K$ is a subset that is closed under addition and under scaling by elements of $\mathcal{O}_K$. Similarly, $\mathcal{J}$ is a fractional ideal if $c\mathcal{J} = \mathcal{I}$ for some ideal $\mathcal{I}$ of $\mathcal{O}_K$ and $c \in \mathcal{O}_K$. Formally, we define an ideal lattice as follows.

Definition 2.20. Let $\mathcal{I}$ be a (possibly fractional) ideal in $\mathcal{O}_K$. The ideal lattice $\sigma(\mathcal{I})$ is the discrete, additive, subset of H that is the image of $\mathcal{I}$ under the canonical embedding σ .

One may restrict computational lattice problems to lattices generated by ideals in some number field K . Since the embeddings provide structure not present on an arbitrary lattice this makes the problems no easier, and any algorithm that solves SVP on an arbitrary lattice would also work on a structured lattice. We present a formal definition for ideal-SVP; definitions of other computational lattice problems are analogous.

Definition 2.21. Let K be a number field together with some norm $\|\cdot\|$ on H and let $\xi = \xi(n) \geq 1$. Then, the Ideal Lattice-SVP $_{\xi}$ problem is as follows: given a (fractional) ideal $\mathcal{I}$ of K , find a short element $x \in \mathcal{I}$ satisfying $\|\sigma(x)\| \leq \lambda_1(\mathcal{I}) \cdot \xi$, where $\lambda_1(\mathcal{I})$ is the minimal norm among elements of $\mathcal{I}$.

Since ideal lattices play a key role in ring based cryptography it is no surprise that understanding the structure of ideals in $\mathcal{O}_K$ is of some importance, so we will quickly run through some important facts and definitions. An integral basis $\{b_1, \dots, b_n\} \subset \mathcal{O}_K$ of $\mathcal{O}_K$ is a $\mathbb{Q}$ -basis of K such that each $x \in \mathcal{O}_K$ can be uniquely represented as $\sum_i b_i z_i$ with each $z_i \in \mathbb{Z}$. Similarly, an integral basis for an ideal $\mathcal{I}$ is any set $\{c_1, \dots, c_n\}$ of elements of $\mathcal{I}$ such that $\mathcal{I} = \{\sum_i c_i z_i \mid z_i \in \mathbb{Z}\}$. For $x \in \mathcal{O}_K$ the ideal generated by x is the smallest ideal containing x , written $\langle x \rangle$, and ideals that can be expressed in this form are called principal ideals. We call $\mathcal{I}$ prime if it is maximal e.g an ideal such that if $\mathcal{I} \subset \mathcal{J} \subset \mathcal{O}_K$ for some proper ideal $\mathcal{J}$ then $\mathcal{I} = \mathcal{J}$. Given two ideals $\mathcal{I}, \mathcal{J}$ we define their sum and product as

$$\mathcal{I} + \mathcal{J} = \{i + j \mid i \in \mathcal{I}, j \in \mathcal{J}\}$$

and

$$\mathcal{I} \cdot \mathcal{J} = \left\{ \sum_{k=1}^z i_k j_k \mid i_k \in \mathcal{I}, j_k \in \mathcal{J}, z \in \mathbb{Z} \right\}.$$

The inverse of an ideal $\mathcal{I}$ is the unique ideal $\mathcal{I}^{-1}$ such that $\mathcal{I} \cdot \mathcal{I}^{-1} = \mathcal{O}_K$. We say that the ideal $\mathcal{I}$ divides $\mathcal{J}$ if there is some ideal $\mathcal{M}$ such that $\mathcal{I} \cdot \mathcal{M} = \mathcal{J}$ and in $\mathcal{O}_K$ it holds that $\mathcal{I} \subset \mathcal{J}$ if and only if $\mathcal{J}$ divides $\mathcal{I}$.

Each ideal has a unique decomposition into a product of prime ideals, and correspondingly we call $\mathcal{I}$ and $\mathcal{J}$ coprime if their prime factorizations' are disjoint, or equivalently if $\mathcal{I} + \mathcal{J} = \mathcal{O}_K$. Denoting the factorization of $\mathcal{I}$ as

$$\mathcal{I} = \prod_i \mathfrak{q}_i^{e_i},$$

we say that $\mathcal{I}$ is ramified if some e_i is greater than 1. Given a further extension L/K , we call a prime K -ideal $\mathfrak{q}$ inert if it remains prime in L . The notion of prime and coprime allows one to define the Chinese Remainder Theorem (CRT) on ideals.

Theorem 2.22 (CRT). *Let $\mathcal{I}$ be an ideal of $\mathcal{O}_K$ which decomposes into a product of coprime ideals as $\mathcal{I} = \prod_{i=1}^r \mathfrak{p}_i$. Then, the ring homomorphism sending $x \in \mathcal{O}_K$ to $(x \bmod \mathfrak{p}_1, \dots, x \bmod \mathfrak{p}_r)$ induces an isomorphism $\mathcal{O}_K/\mathcal{I} \rightarrow \bigoplus_i \mathcal{O}_K/\mathfrak{p}_i$.*

Since ideals generate lattices, it is intuitive to define the notion of a dual ideal.

Definition 2.23. Let $\mathcal{I}$ be an ideal of K . Then, the dual ideal $\mathcal{I}^\vee$ is the unique ideal

$$\mathcal{I}^\vee = \{x \in K \mid \text{Tr}_{K/\mathbb{Q}}(x \cdot \mathcal{I}) \subset \mathbb{Z}\}.$$

It is easy to check that the dual ideal is indeed a (possibly fractional) ideal of K , and that the ideal lattice generated by $\mathcal{I}^\vee$ embeds as the complex conjugate of the dual lattice of the embedding of $\mathcal{I}$. A dual ideal of great importance is the codifferent ideal $\mathcal{O}_K^\vee$, which is the name given to the dual of $\mathcal{O}_K$ itself, and it can be seen that for any ideal $\mathcal{I}$ we have $\mathcal{I}^\vee = \mathcal{I}^{-1} \cdot \mathcal{O}_K^\vee$.

2.5 Ring Learning With Errors

We begin with the dual form definition of Ring LWE from [LPR10], because it will help us with the formalizing of CLWE later. Afterwards, we introduce the conceptually simpler primal form, which is substantially more common in practical applications. We start, as usual, with a distribution. We will use the notation common to RLWE papers that $\mathcal{O}_{K_q} := \mathcal{O}_K / \langle q \rangle$, and more generally $\mathcal{I}_q := \mathcal{I} / q\mathcal{I}$ for any ideal $\mathcal{I}$ and integer q .

Definition 2.24 (Ring-LWE Distribution). Given a secret $s \in \mathcal{O}_{K_q}^\vee$ and an error distribution ψ over $K_{\mathbb{R}}$, a sample from the ring LWE distribution $A_{s,\psi}$ is obtained by sampling $a \leftarrow \mathcal{O}_{K_q}$ uniformly at random, sampling $e \leftarrow \psi$, and outputting $(a, (a \cdot s)/q + e \bmod \mathcal{O}_K^\vee) \in \mathcal{O}_{K_q} \times K_{\mathbb{R}}/\mathcal{O}_K^\vee$.

As with standard LWE the RLWE distribution enables a search problem and a decision problem, which are defined as follows:

Definition 2.25. Let $q \geq 2$ be an integer and let Ψ be a family of error distributions on $K_{\mathbb{R}}$. Then, the search version of the RLWE problem, denoted $\text{RLWE}_{q,\Psi}$, is to compute s given an arbitrary number of samples from $A_{s,\psi}$ for some $s \in \mathcal{O}_{K_q}^\vee, \psi \in \Psi$.

The (average-case) decision version of the RLWE, denoted $\text{D-RLWE}_{q,\Upsilon}$, is to distinguish with non-negligible advantage between the uniform distribution on $\mathcal{O}_{K_q} \times K_{\mathbb{R}}/\mathcal{O}_K^\vee$ and the RLWE distribution $A_{s,\psi}$, where s is chosen uniformly at random from $\mathcal{O}_{K_q}^\vee$ and ψ is chosen according to some arbitrary distribution Υ on Ψ .

Before discussing the hardness of these problems let us consider the choice in the construction of RLWE as to where we draw a and s from, which turn out to be motivated by the security reduction in [LPR10]. As part of a security proof they provide an ideal-BDD to RLWE reduction in which the element s corresponds to the closest vector in the lattice $\mathcal{I}^\vee$ for some instance of the BDD problem and a is drawn from a Gaussian sample on the lattice $\mathcal{I}$. This ideal is then cleared by an isomorphism to leave us with uniform $s \in \mathcal{O}_{K_q}^\vee, a \in \mathcal{O}_{K_q}$, explaining the construction of the RLWE distribution. Intuitively, one might expect that the ring based

LWE problems are approximately as hard as the equivalent standard LWE problems, since the underlying structure of the problems are similar. In [LPR10], Lyubashevsky et al. present an efficient quantum reduction from ideal-SIVP (or ideal-SVP) to search RLWE, as well as internal reductions similar to those of the LWE problem such as search-decision and worst-case to average-case. We present their main theorem and a summary of the key steps, but not the full proof. In the following theorem a sample from the distribution Υ_α is taken by sampling $x_1, \dots, x_{n/2}$ from $\Gamma(2, 1)$ and outputting a sample from the elliptical Gaussian $D_{\mathbf{r}}$, where $\mathbf{r} = (r_1, \dots, r_n)$ for $r_i^2 = r_{i+n/2}^2 = \alpha^2(1 + \sqrt{n}x_i)$.

Theorem 2.26. *Let K be the m^{th} cyclotomic number field with ring of integers $\mathcal{O}_K$, $n = \phi(m)$, and $\alpha < \sqrt{\frac{\log n}{n}}$. Furthermore, let $q = q(n)$ be a $\text{poly}(n)$ bounded prime number satisfying $q \equiv 1 \pmod{m}$ and $\alpha q \geq \omega(\sqrt{\log n})$. Then, there exists a polynomial time quantum reduction from $\tilde{O}(\frac{\sqrt{n}}{\alpha})$ -approximate SVP (or SIVP) on ideal lattices in K to $D\text{-RLWE}_{q, \Upsilon_\alpha}$.*

The reduction to search RLWE drops the hypothesis that K is a cyclotomic number field and the corresponding restriction that $q \equiv 1 \pmod{m}$, hypotheses that are necessary only for the search-decision reduction. The reduction from ideal-SVP to search RLWE follows the skeleton of Regev's reduction for classic LWE, whereas the search-decision and worst-case to average-case reductions utilize more complex machinery involving reducing modulo the ideal factors $\mathfrak{q}_1, \dots, \mathfrak{q}_n$ of $\langle q \rangle$, guessing the value of s modulo each of the $\mathfrak{q}_i$, then recovering s using the CRT.

2.5.1 Primal Form, Normal Form, and RLWE as Structured LWE

In order to present RLWE as a structured variant of LWE we must first discretize the samples. In [LPR13] they show that the search and decision problems are still hard in so-called 'primal' form, where the secret and error no longer come from the dual ideal but instead come from the ring of integers itself. Hence, we can discretize the distribution without loss of security.

Definition 2.27 (Primal RLWE Distribution). For a secret $s \in \mathcal{O}_{K_q}$ and an error distribution χ over $\mathcal{O}_{K_q}$ the primal form of the RLWE distribution, again denoted $A_{s, \chi}$, is sampled from by

choosing $a \leftarrow \mathcal{O}_{K_q}$ uniformly at random, sampling $e \leftarrow \chi$, and outputting $(a, b := a \cdot s + e)$.

Remark. To avoid confusion we will denote the discrete error distribution by χ when $A_{\mathbf{s},\chi}$ is the primal form of the RLWE distribution and use ψ to denote the continuous distribution when $A_{\mathbf{s},\psi}$ is in dual form.

The definition of the search and decision problem remain unchanged, except that the uniform distribution in the decision problem is now over the ambient space $\mathcal{O}_{K_q} \times \mathcal{O}_{K_q}$. In order to maintain security in primal form one must ‘tweak’ the error distribution χ (see [CP16]) to compensate for the lower density of the non-dual ring of integers $\mathcal{O}_K$, but up to this minor change the variants have equivalent security. As in the case of classical LWE, [LPR13] shows that the RLWE problem in primal form is still hard in normal form, where the secret s is sampled from the error distribution χ rather than uniformly at random from $\mathcal{O}_{K_q}$.

We are now ready to formalize the notion of RLWE as a structured instance of LWE with number of samples equal to the dimension n . Consider instantiating the normal form of primal RLWE in the two-power cyclotomic field $K = \mathbb{Q}(\zeta_m)$ for m a power-of-two, and let n denote $m/2$, the dimension of K . Then, the ring of integers satisfies $\mathcal{O}_K = \mathbb{Z}(\zeta_m)$ and $\mathcal{O}_{K_q} \cong R_q := \mathbb{Z}_q[x]/(x^n + 1)$. Let χ denote the error distribution over R_q which is sampled from by sampling each coefficient independently from the 1-dimensional discrete Gaussian with parameter r over $\mathbb{Z}_q$, and for $a \in R_q$ let $\mathbf{a}$ denote its coefficient vector over $\mathbb{Z}_q$. Then, an RLWE sample in R_q takes the form $(a, b := a \cdot s + e)$, and it is straightforward to show that this equation is equivalent to the matrix-vector equation $(\mathbf{A}, \mathbf{b} := \mathbf{A}\mathbf{s} + \mathbf{e})$ over $\mathbb{Z}_q$, where $\mathbf{A} = \text{rot}(a)$. Explicitly, this equation takes the form:

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

with $a_i \in \mathbb{Z}_q$ denoting the i^{th} coefficient of the polynomial a . Now it is clear that this is almost an LWE instance with secret $\mathbf{s}$ and error distribution D_r , the difference being that the n

vectors $\mathbf{a}_i$ are no longer independent and uniformly random but instead come with a rotational structure. Remarkably, in common applications where one has only a couple of samples from the RLWE distribution, such as in NewHope [ADPS16], there are no known methods that exploit this cyclical structure to solve the RLWE problem. Instead the matrix $\text{rot}(a)$ is treated as unstructured and a single sample from the RLWE distribution is treated and solved as n independent samples from the LWE distribution.

2.6 Module LWE

Module LWE, introduced in [LS15], generalizes ring and classical LWE by instantiating the problems in a module over a ring. Due to its conceptual simplicity we will present only the primal form, but MLWE permits a dual form as in RLWE which one can deduce from its primal variant. We will need the definition of a module.

Definition 2.28. Let $M \subset K^d$. We call M an $\mathcal{O}_K$ -module if it is closed under addition and under scaling by elements of $\mathcal{O}_K$.

The module used in MLWE is the full $\mathcal{O}_K$ -module, $M := \mathcal{O}_K^d$. One defines the inner product $\langle \cdot, \cdot \rangle$ on M as the usual vector space inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^d x_i y_i$. Once again, we will begin our definition of the MLWE problems with a distribution.

Definition 2.29. Let K be a number field, $q \geq 2$ and $d \geq 1$ be integers, and let $M = \mathcal{O}_{Kq}^d$. Then for an error distribution χ over $\mathcal{O}_{Kq}$ and a secret $\mathbf{s} \in M$ a sample from the MLWE distribution $A_{\mathbf{s}, \chi}$ is obtained by sampling $\mathbf{a} \leftarrow M$ uniformly at random, sampling $e \leftarrow \chi$, and outputting $(\mathbf{a}, b = \langle \mathbf{a}, \mathbf{s} \rangle + e)$.

Remark. From the MLWE distribution one can recover the RLWE distribution in dimension n by setting $d = 1$ or the LWE distribution in dimension d by setting $n = 1$. Thus, the MLWE distribution and associated problems generalize those of LWE and RLWE.

Given this distribution we can define the search and decision problems as usual: find $\mathbf{s}$, or distinguish MLWE samples from uniform samples over $M \times \mathcal{O}_{Kq}$. [LS15] provides all the usual

LWE security reductions for MLWE: search to average-case decision together with a reduction from certain structured lattice problems to the search problem with Gaussian error. This time, the appropriate short vector problems are over module lattices generated by applying the canonical embedding σ of K coordinatewise on appropriate submodules of M . Since M has dimension d over $\mathcal{O}_{K_q}$ and K has dimension n this can be seen as a structured lattice problem over $\mathbb{R}^N$ for $N := nd$, the total dimension of the MLWE instance over $\mathbb{Z}_q$.

2.6.1 MLWE as Structured LWE

Given d samples from the MLWE distribution $A_{\mathbf{s},\chi}$ one can once again recover a structured variant of the LWE problem. As before, choose K as the m^{th} cyclotomic number field with $m = 2n$ a power-of-2 such that $\mathcal{O}_{K_q} \cong R_q$ and collate the d samples into a matrix $\mathbf{A} \in R_q^{d \times d}$, error vector $\mathbf{e}$, and corresponding vector $\mathbf{b}$. Then, the d samples from the MLWE distribution can be written as $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$, where each vector lies in R_q^d . Over R_q , this forms the matrix-vector equation

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

To transform this into a matrix-vector equation in dimension $N = nd$ over $\mathbb{Z}_q$ one applies the $\text{rot}(a_{i,j})$ operation coordinatewise on $\mathbf{A}$ and expresses each $s_i \in R_q$ in its coefficient vector over $\mathbb{Z}_q$. This results in the now familiar form $(\mathbf{A}, \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e})$ over $\mathbb{Z}_q$, where $\mathbf{A}$ now has the blockwise rotational structure of Fig. 1.1.

2.7 Structured LWE in Cryptography

We would like to briefly compare the relative merits of ring and module LWE, since neither should be viewed as strictly better than the other. To provide a basis for our comparison

we present an example of LWE based public key cryptography, first presented in [LPR10]. Extensions of this scheme to a KEM are used in several remaining NIST process candidates, including NewHope and Kyber, and this scheme can be considered as essentially the canonical example of LWE variants in public key cryptography. We present the ring variant due to its conceptual simplicity, but it is straightforward to construct corresponding schemes using classical or module based LWE as well.

The scheme is parameterized by an integer modulus q , dimension n , and error distribution χ over R_q . We represent the ring R_q as $\frac{\mathbb{Z}_q[X]}{x^n + 1}$ and define χ as the distribution over R_q obtained by sampling each coefficient of a polynomial independently from the one dimensional discrete Gaussian over $\mathbb{Z}_q$ of standard deviation r . The scheme proceeds as follows.

- Alice samples a secret key $s \leftarrow \chi$ and publishes as a public key an RLWE sample $(a, b) = (a, a \cdot s + e) \in R_q \times R_q$, where a is uniformly random and $e \leftarrow \chi$.
- Bob encrypts a message $m \in \{0, 1\}^n$ as $(c_1, c_2) = (a \cdot t + e_1, b \cdot t + e_2 + \lfloor \frac{q}{2} \rfloor \cdot m)$, where e_1, e_2, t are sampled independently from χ .
- Alice decrypts using s by computing $c := c_2 - c_1 \cdot s = \lfloor \frac{q}{2} \rfloor \cdot m + e \cdot t - s \cdot e_1 + e_2$.

Alice then recovers the message m by decoding: if the i^{th} coefficient of c is closer to 0 than $\lfloor q/2 \rfloor$ then Alice assumes the i^{th} coordinate of m was 0, otherwise she assumes it was 1. Since e, e_1, e_2, t and s were sampled independently from the error distribution χ , each coefficient of the combined error term $e \cdot t - s \cdot e_1 + e_2$ is smaller than $\lfloor \frac{q}{4} \rfloor$ with high probability and hence this decoding almost always outputs the correct message m . Although the schemes formal security proof may be found in [LPR13], the main idea is simply that b, c_1 , and c_2 leak no information because they are RLWE samples and thus appear pseudorandom to an onlooker by the hardness of the RLWE decision problem.

To convert this scheme to an MLWE or LWE based scheme, recall that by using the matrix representation of multiplication in R_q all these LWE variants can be considered as structured instantiations of the equation $\mathbf{b} = \mathbf{A}s + \mathbf{e}$. Replacing the polynomial a in the scheme by a

matrix $\mathbf{A}$ with structure that is uniformly random, derived from $\text{rot}(a)$, or derived from blocks of $\text{rot}(a_{i,j})$ bases the scheme on LWE, RLWE, and MLWE respectively.

2.7.1 Comparing Ring and Module LWE

Now that we have a simple model for public key schemes based on LWE, we are ready to compare the relative advantages of the two structured variants of LWE. Since in practice the difficulty of the LWE problem is currently treated as independent of the matrix structure we will assume in our comparison that the total dimension N , defined by the dimension of the square matrix $\mathbf{A}$, is fixed over the choice of LWE assumption and thus base our comparison on schemes in the same dimension where possible.

First, we consider the storage requirements the scheme has of its ambient space. The public key consists of the matrix $\mathbf{A}$ and the element b . Since over all variants $\mathbf{b}$ can be stored as an unstructured vector over $\mathbb{Z}_q$ in dimension N , the cost of storing this key is asymptotically dominated by that of storing $\mathbf{A}$. In the ring case, $\mathbf{A}$ can be stored as an element of R_q in dimension $N = n$, whereas in the module case $\mathbf{A}$ must be stored as d^2 entries in R_q of dimension n . Here the storage advantages of RLWE are clear: an RLWE key can store the public key as two vectors of length N over $\mathbb{Z}_q$, but the equivalent MLWE scheme stores the matrix as d^2 vectors of length n , requiring a total of $nd^2 = Nd$ elements of $\mathbb{Z}_q$, as well as the single vector $\mathbf{b}$ of length N .

Second, we consider the efficiency of the encryption and decryption phases. Both consist solely of matrix-vector operations such as $\mathbf{A}\mathbf{s}$ and some vector addition. Since there are only few additions it is clear that the asymptotic cost of this operation will be dominated by the speed of multiplication within the ambient space. In both cases the method of multiplication will be to apply the CRT decomposition from [LPR13], which computes a multiplication in R_q in $O(n \log n)$ operations over $\mathbb{Z}_q$. In the ring case $n = N$ and so the asymptotic complexity is just $O(N \log N)$. In the module case the dimension of R_q is n and the $d \times d$ matrix over R_q results in d^2 instances of the CRT in dimension n , giving a total complexity of $O(Nd \log n)$ over $\mathbb{Z}_q$.

This is typically a little slower since $d \log n$ grows faster than $\log nd$ as n increases. Once again, RLWE comes out on top.

However, this additional efficiency does come at a cost in flexibility. The dimension of the ring R_q is chosen so that n is a power-of-2, and suitable cryptographic dimensions for LWE based schemes often lie somewhere between 512 or 1024. This restriction on n means that RLWE based schemes are restricted to exactly these dimensions. On the other hand, module schemes have flexible control over the parameter d , and setting for instance $d = 3$ and $n = 256$ allows for a total dimension of 768, which may be preferable depending on the targeted level of security. Of course, a ring based scheme can achieve that security level or more in dimension 1024, but the overkill in dimension results in meaningfully slower operations. There is another subtle benefit in the module case here, that one may change dimension by fixing the ring R_q and increasing the rank d , whereas a ring scheme is forced to increase the dimension of R_q . This allows a module scheme to specialize its R_q multiplication algorithm to a fixed choice of n , leading to practical, although not asymptotic, speed ups. Conversely, a ring based scheme must either have a less efficient but more general algorithm for variable choices of n or develop individual specialized algorithms for each dimension. This is exactly the methodology used by Kyber, where there is a specialized multiplication algorithm for dimension $n = 256$ and then security levels are altered by increasing d [BDK⁺18].

Finally we must consider the effect of the choice of ambient space on the concrete security of the scheme. Here, it seems that MLWE is advantaged, although this is a relatively open area of study. Since the security of each scheme is based on the difficulty of differentiating samples from the respective LWE distributions from random, the hardness of each scheme corresponds to that of the decision LWE problem in the appropriate ambient space. It seems that the algebraic structure of the RLWE problem may lead to some future vulnerabilities (see e.g. [Ber14], [CDPR16], [BCLvV17]), and that these attacks would likely be less effective on MLWE. However, this is a poorly understood problem and for the moment the typical method for analyzing security is the CoreSVP technique of NewHope [ADPS16], which treats both cases as the unstructured LWE problem and so leads to the same security estimates.

We've now concluded our introduction to the LWE problem and its ring and module variants. We will return in more detail to parts of specific importance in the main body of the thesis, but for the time being move on to a description of the main algebraic object of this work, cyclic algebras.

2.8 Cyclic Algebras

Typically, an introduction to cyclic algebras begins with the underpinnings from algebraic number theory. Fortunately, we have already introduced the necessary concepts and so we proceed immediately to the definition of a cyclic algebra.

Definition 2.30. Let K be a number field of degree n and let L be a cyclic Galois extension of K of degree d with $\text{Gal}(L/K) = \langle \theta \rangle$. For a non-zero $\gamma \in K$ we define the resulting cyclic algebra

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

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 $u^d = \gamma$. From these conditions it follows that multiplication in $\mathcal{A}$ is non-commutative. We call d the degree of the algebra $\mathcal{A}$. We call $\mathcal{A}$ a division algebra if every element $a \in \mathcal{A}$ has an inverse $a^{-1} \in \mathcal{A}$ such that $aa^{-1} = 1$. The dimension of $\mathcal{A}$ is defined as its dimension $N := nd^2$ over $\mathbb{Q}$.




Figure 2.3: Construction of a Cyclic Algebra.

Remark. There are instances in the literature where the direct sum is taken with u on the right of L and the multiplicative relation is replaced with the corresponding relation $ux = \theta(x)u$. This leads to an equivalent but distinct algebra $\mathcal{A}'$ with essentially the same properties. In this thesis we will always place the indexing element u on the left.

Since θ fixes K , the center of the cyclic algebra is precisely K . We will sometimes refer to L as the maximal subfield of $\mathcal{A}$. Oftentimes the condition $\gamma \in K$ is replaced by the stronger condition $\gamma \in \mathcal{O}_K$, and we will enforce this condition in our work unless otherwise specified. Note that the division property does not hold for arbitrary γ and that cyclic algebras with the division property, which we refer to as cyclic division algebras, are not always easy to construct.

2.8.1 Matrix Representation

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

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

We call this mapping a left regular representation of $\mathcal{A}$ because it holds for any $a, b \in \mathcal{A}$ that $\phi(a) \cdot \mathbf{b} = \text{Vec}(ab)$, and that $\phi(ab) = \phi(a) \cdot \phi(b)$, from which it follows that ϕ is an embedding. Since ϕ is homomorphic it follows that if $\mathcal{A}$ is a division algebra then $\phi(a)$ is an invertible matrix for each $a \in \mathcal{A}$. Since θ can be extended to a well-defined automorphism of $L_{\mathbb{R}}$ we abuse notation and extend this map to a homomorphism $\phi : \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}} \rightarrow M_{d \times d}(L_{\mathbb{R}})$.

2.8.2 Orders and the Natural Order

The definition of the RLWE distribution is not set in the number field K itself, rather in its ring of integers $\mathcal{O}_K$, typically in polynomial form R_q . Now that we've introduced cyclic algebras,

we need to find our analog of R_q .

Definition 2.31. 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 $\mathcal{A}$ which contains the identity. We call Λ maximal if there is no $\mathbb{Z}$ -order Γ such that $\Lambda \subsetneq \Gamma \subsetneq \mathcal{A}$. Here, $\Lambda \cdot \mathbb{Q} = \{\sum_{i=1}^m a_i q_i : a_i \in \Lambda, q_i \in \mathbb{Q}, m \in \mathbb{Z}_{\geq 1}\}$.

As we are only concerned with $\mathbb{Z}$ -orders in this thesis we will just refer to them as orders. Nonetheless, it is possible to generalize the notion of orders to e.g. $\mathcal{O}_K$ -orders by enforcing that Λ is an $\mathcal{O}_K$ -module instead of a $\mathbb{Z}$ -order and that $\Lambda \cdot K = \mathcal{A}$. Some results we use on $\mathbb{Z}$ -orders are typically proven for the more general $\mathcal{O}_K$ -orders, but since being a $\mathbb{Z}$ -order is a strictly stronger condition we only use those.

Remark. 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.

Definition 2.32. Given an algebra $\mathcal{A} = (L/K, \theta, \gamma)$ with $\gamma \in \mathcal{O}_K$, the *natural order* Λ is defined as $\Lambda := \bigoplus_{i=0}^{d-1} u^i \mathcal{O}_L$.

Unlike in the case of $\mathcal{O}_K$, this order is not necessarily maximal. Note that in order for Λ to be closed under multiplication the element γ must lie in $\mathcal{O}_K$, which is the reason this condition is often enforced.

2.8.3 Existence and non-Norm Elements

It is not a priori obvious whether well-defined cyclic division algebras actually exist. As observed earlier, the existence of γ enforcing the division condition is a key component in constructing such objects. Fortunately, it is sufficient (and necessary) for γ to satisfy the so called ‘non-norm condition’. Observe that since $\gamma \in K$ it is always true that $\gamma^d \in N_{L/K}(L^\times)$, where $L^\times$ denotes L without 0. The non-norm condition is satisfied by γ if γ^d is the lowest power of γ that appears in $N_{L/K}(L^\times)$, in which case the cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ is a cyclic division algebra.

Theorem 2.33. *Let $\mathcal{A} := (L/K, \theta, \gamma)$ be a cyclic algebra with $[L : K] = d$. Then, $\mathcal{A}$ is a division algebra if and only if the smallest non-zero power of γ that appears in $N_{L/K}(L^\times)$ is γ^d .*

This result is well known and a proof may be found in many coding papers on cyclic algebras, for example in [VHLR09], and we follow the usual convention by referring to such a γ as a non-norm element. Recall that in order to keep the natural order well-defined we will still be interested in γ being an integral element of K , so we now have two substantial yet distinct restrictions on γ . Consequentially, when we require algebras satisfying certain conditions we will be careful to prove that such algebras actually exist.

2.8.4 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, although we will primarily be interested in only one. 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. In the following we restrict our interest to two-sided ideals and refer to them merely as ideals, although some results and definitions also hold for one-sided ideals as well. 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, that $\mathcal{I}$ is a fractional ideal of Λ if $c\mathcal{I} = \mathcal{J}$ and $\mathcal{I}$ for a two-sided ideal $\mathcal{J}$ and some central $c \in \mathcal{O}_K$.

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 [Rei75] states that the multiplicative 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 invertible two-sided case (e.g. $\mathcal{I} \cdot \mathcal{I}^{-1} = \mathcal{I}^{-1} \cdot \mathcal{I} = \Lambda$).

Since in a natural order we do not have a well-behaved ideal group, we will rely on the exposition of Section 3 of [OAS12] to understand their structure. In particular, we will use the fact that for a two-sided ideal $\mathcal{I} \subset \Lambda$ the intersection $\mathcal{I} \cap \mathcal{O}_K$ is an ideal of $\mathcal{O}_K$ which one can compute efficiently. 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 $\mathcal{O}_K$, but it is not in general injective since several ideals of $\mathcal{A}$ may have the same intersection with $\mathcal{O}_K$. Since the ideals of Λ do not necessarily form a finitely generated abelian group we define two ideals $\mathcal{I}, \mathcal{J}$ of Λ to be coprime if $\mathcal{I} + \mathcal{J} = \Lambda$.

We will also use the following decomposition, Lemma 4 of [OAS12], which we describe as a ‘CRT-like’ decomposition throughout this thesis. We will refer to certain details of the proof later in this work, and so present the key steps here.

Lemma 2.34. *Let $\mathcal{I}$ be an ideal of $\mathcal{O}_K$ which factorizes into a product of prime ideals as $\mathcal{I} = \mathfrak{q}_1^{e_1} \dots \mathfrak{q}_t^{e_t}$ in $\mathcal{O}_K$. Then,*

$$\Lambda/\mathcal{I} \cdot \Lambda \cong \mathcal{R}_1 \times \dots \times \mathcal{R}_t$$

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

Proof Outline. Since the $\mathfrak{q}_i$ are prime as $\mathcal{O}_K$ ideals they are coprime ideals in $\mathcal{O}_L$, so we may apply the usual CRT to construct the isomorphism $\mathcal{O}_L/\mathcal{I}\mathcal{O}_L \cong \mathcal{O}_L/\mathfrak{q}_1^{e_1}\mathcal{O}_L \times \dots \times \mathcal{O}_L/\mathcal{O}_L\mathfrak{q}_t^{e_t}$. From the discussion of surjection of ideals it follows that $\Lambda/\mathcal{I}\Lambda \cong \bigoplus_{i=0}^{d-1} u^i(\mathcal{O}_L/\mathcal{I}\mathcal{O}_L)$, and so applying the CRT operation coordinatewise over the u coordinates of $\Lambda/\mathcal{I} \cdot \Lambda$ gives the desired result. $\square$

Typically we will apply this in the case where $\mathcal{I} = \langle q \rangle$, the principal ideal generated by some prime integer q .

2.8.5 Some Useful Ideals

As in the case of RLWE we will be interested in the notion of dual ideals and principal ideals. Since the structure of a cyclic algebra requires a small change in their definitions we revisit them explicitly. For an order Λ we define the codifferent ideal

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

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

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

This definition is suitably two-sided because the matrix trace satisfies $\text{Trace}(AB) = \text{Trace}(BA)$. Note that the codifferent ideal and a general dual ideal may be fractional ideals rather than full ideals, and as in the ring case 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 simply define $\langle t \rangle = t \cdot \Lambda$, the set of elements of Λ divisible by t . However, for a general t that does not lie in the center of Λ , e.g. those elements that do not commute with all of $\mathcal{A}$, we need a slightly more complex definition because $t \cdot \Lambda$ is not closed under left multiplication. Instead, we use the ideal $\langle t \rangle := \Lambda \cdot t \cdot \Lambda$, written more explicitly as

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

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

2.8.6 Orders and Ideals as Integer Lattices

Any order Λ of a cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ has dimension $N := nd^2$ over $\mathbb{Z}$ and thus generates a $\mathbb{Z}$ -lattice of dimension N . We will utilize the following representation of these lattices, which extends immediately to ideals of orders as well. Given an element $x = \bigoplus_{i=0}^{d-1} u^i x_i \in \Lambda$, we can consider x as a vector over H_L of dimension d by defining the map $\sigma_{\mathcal{A}}(x) := (\sigma_L(x_0), \sigma_L(x_1), \dots, \sigma_L(x_{d-1}))^1$. Then, the collection $\sigma_{\mathcal{A}}(\Lambda)$ forms an integer lattice of dimension nd^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 $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$. Though this representation is conceptually simple we remark that when considering norms of lattice elements it has some drawbacks in the case where $|\sigma_i(\gamma)| \neq 1$ for some i ; we will choose γ carefully in our constructions to remove this issue². Similarly, an ideal $\mathcal{I}$ of Λ generates a lattice via $\sigma_{\mathcal{A}}(\mathcal{I})$ and we refer to these lattices as ideal lattices in Λ , or sometimes order-ideal lattices to prevent confusion with ideal lattices derived from number fields.

2.8.7 Gaussian Distributions Over Cyclic Algebras

As with other forms of LWE we will need to sample Gaussian distributions over our ambient space in appropriate norms. In the case of RLWE the continuous Gaussians are sampled in $K_{\mathbb{R}} \cong H$. Since a cyclic algebra $\mathcal{A}$ can be viewed as a d -dimensional algebra over L we make use of the visualization from the previous subsection and sample our error distributions over $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$, which has the same structure as a vector space as H_L^d . For simplicity, we restrict ourselves to the case where $|\sigma_i(\gamma)| = 1$ 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. For example, if γ is ‘large’ then later coordinates should have larger weight since after multiplication they will be hit by γ more often. Additionally, the relative size of the L -embeddings of γ would

¹Unlike in the ring case this is not a homomorphism as it is not multiplicative.

²In cases where some $|\sigma_i(\gamma)| \neq 1$ it is customary to generate lattices by vectorizing the matrix representation $\phi(a)$ and applying L -extensions of K -embeddings, but we will not require this technique.

need to be taken into account. These complications are substantial, hence we find it easier to just restrict γ .

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^d of dimension N used in [LS15], e.g.

$$|x| := \|x\|_2 = \|(\sigma_L(x_0), \sigma_L(x_1), \dots, \sigma_L(x_{d-1}))\|_2$$

for $x = x_0 + ux_1 + \dots + u^{d-1}x_{d-1} \in \mathcal{A}$. It is straightforward to check that this is indeed a norm in the case where $|\sigma_i(\gamma)| = 1$ for each i , since γ is fixed under θ and multiplying by γ does not change the norm of an entry of σ_L . It is also clear that this norm extends to any $y \in \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ by considering the extension of the norm on L to $L_{\mathbb{R}}$ in each coordinate. 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 $L_{\mathbb{R}}^d$ or $\mathcal{O}_L^d$ and then reconstructing an element of $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ in the natural manner.

Chapter 3

Learning With Errors over Cyclic Algebras

But here's some advice, boy. Don't put your trust in revolutions. They always come around again. That's why they're called revolutions.

Sam Vimes, *Night Watch* by Terry

Pratchett

In this chapter we will introduce the central contribution of this thesis, the construction of the Learning with Errors problem over cyclic algebras, which we call Cyclic LWE or just CLWE. We will begin by defining a distribution in the style of the dual variant of the Ring LWE assumption. Before providing a security reduction to lattice problems we will provide a definition for standard lattice problems, such as SVP and BDD, over two-sided ideals of the natural order of a cyclic algebra, and we must adapt some technical lemmas about $\mathcal{O}_K$ -ideals used in the RLWE security proof of [LPR10] to order-ideals. Following this, the main part of this section will be the adaptation of the standard security reductions for an LWE problem for the CLWE case: a reduction from Search CLWE to structured lattice short vector problems and a reduction from average-case Decision CLWE to Search CLWE. However, our search-decision reduction only

holds for a variant of Search CLWE in which the secret s is restricted to certain large subsets of $\Lambda_q^\vee$. Afterwards we will also discuss why the direct-to-decision reduction for RLWE from [PRSD17], which is also valid for MLWE, does not adapt naturally to the case of CLWE.

3.1 Definition of CLWE

We begin with the definition of the CLWE problems in order to provide motivation for the rest of the section. As in the ring case we will use the notation $\mathcal{I}_q$ for the quotient $\mathcal{I}/q \cdot \mathcal{I}$ of any (two-sided) ideal $\mathcal{I}$ of the natural order Λ and integer q .

Definition 3.1. Let L/K be a Galois extension of number fields of dimension $[L : K] = d$, $[K : \mathbb{Q}] = n$ with cyclic Galois group generated by $\theta : L \rightarrow L$. Let $\mathcal{A} := (L/K, \theta, \gamma)$ be the resulting cyclic algebra with center K and invariant u satisfying $u^d = \gamma \in \mathcal{O}_K$. Let Λ be the natural order of $\mathcal{A}$. For an error distribution ψ over $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$, 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, sampling $e \leftarrow \psi$, and outputting $(a, b) = (a, (a \cdot s)/q + e \bmod \Lambda^\vee) \in \Lambda_q \times (\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}})/\Lambda^\vee$.

Remark. We have applied the restriction $\gamma \in \mathcal{O}_K$ in order to maintain the closure of the natural order and hence keep our distribution well-defined. The other restriction on γ noted in the previous chapter, that $|\sigma_i(\gamma)| = 1$ for each embedding σ_i of L , will not be needed until we discuss security properties.

As is the case with other LWE assumptions, the associated CLWE problem will come in search and decision variants.

Definition 3.2. 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 $\text{CLWE}_{q,s,\psi}$, 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. However, we remark that our security reduction will permit only a

polynomial number of samples to an oracle for the Search CLWE problem, and our toy model of a public key encryption scheme will follow the example of NewHope and restrict to at most two samples for a single key s .

Definition 3.3. Let Υ be some distribution on a family of error distributions Σ_α over $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ and let U_Λ denote the uniform distribution on $\Lambda_q \times (\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}) / \Lambda^\vee$. Then, the decision CLWE problem, written $\text{D-CLWE}_{q,\Upsilon}$, is on input a collection of independent samples from either $\Pi_{q,s,\psi}$ for a random choice of $(s, \psi) \leftarrow U(\Lambda_q^\vee) \times \Upsilon$ or from U_Λ to decide which is the case with non-negligible advantage.

Remark. The CLWE distribution and associated problems are in fact a generalization of their ring analogs, since in the case where $d = 1$ the extension L/K is the trivial extension $L = K$ and it is easy to check that the CLWE distribution is precisely the RLWE distribution over K with secret s and error distribution ψ .

When we refer to the CLWE assumption we mean that we are assuming that these problems are computationally infeasible i.e. that there are no polynomial time algorithms that solve them. As is typically the case for LWE cryptography, the decision problem will form the basis of our argument for cryptographic security because it allows us to argue that samples from the CLWE distribution are pseudorandom. If the samples are pseudorandom, a cryptosystem based on some collection of CLWE samples will provide attackers with no information on either the secret key or the encrypted message. If only the search problem were hard then the guarantee is that an attacker would not be able to recover the secret key from some collection of samples, but they may be able to distinguish some information about the message without recovering said key. That is, hardness of the decision problem provides a simple scheme based on the CLWE distribution with IND-CPA security. If required, a standard transform such as the Fujisaki-Okamoto transform [FO99] can then be applied to give the stronger notion of CCA security. That said, even if only the search variant were hard one can still recover an IND-CPA secure KEM using more intricate constructions involving hardcore bits, for example using the techniques of [SSTX09].

3.2 Lattice Problems Over Cyclic Algebras

Since we intend to reduce the CLWE problem to certain structured lattice problems, we must define the specific lattices on which we set these computational problems. The lattices will be generated by embedding ideals of the natural order Λ and so we refer to the problems as order-ideal lattice problems. We will embed an ideal $\mathcal{I}$ into a vector space of dimension nd^2 using the map $\sigma_{\mathcal{A}} : \mathcal{A} \rightarrow H_L^d$ from Section 2.8.6, defined

$$\sigma_{\mathcal{A}}\left(\sum_{i=0}^{d-1} u^i x_i\right) := (\sigma_L(x_0), \dots, \sigma_L(x_{d-1})).$$

We will also use the extension of $\sigma_{\mathcal{A}}$ to $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ to construct an ambient space for sampling error and defining the bounded distance decoding problem. Restricting to the case where $|\sigma_i(\gamma)| = 1$ for each i , this allows us to define the ℓ_2 norm on $\mathcal{A}$ by setting $|x| := \|x\|_2 = \|\sigma_{\mathcal{A}}(x)\|_2$. This is the corresponding norm when considering the comparison to MLWE and thus is the norm we will pick for our lattice problems, since it is fixed under the algebra to module mapping $\text{Mod}(\cdot) : \mathcal{A} \rightarrow L^d$, $\text{Mod}(x) = (x_0, \dots, x_{d-1})$. Using the Mod mapping we could also define other standard norms such as the infinity norm, or even module norms from [LS15] such as the $\ell_{2,\infty}$ and $\ell_{\infty,2}$ norms. Since the last two are non-standard, we will be careful to define them when we use them. We are now ready to define order-ideal lattice problems. We define the problems only on two-sided ideals, so for brevity we refer to them merely as ideals.

Definition 3.4. Let $\mathcal{A}$ be a cyclic algebra and let $\mathcal{I}$ be some (possibly fractional) ideal of the natural order Λ . Then, for an approximation factor $\xi = \xi(n) \geq 1$ the (approximate) Order-Ideal Shortest Vector Problem, denoted $\mathcal{A}\text{-SVP}_{\xi}$, is to find a non-zero element $a \in \mathcal{I}$ such that $|a| \leq \xi \cdot \lambda_1(\mathcal{I})$, where as usual $\lambda_1(\mathcal{I})$ denotes the minimal length of elements of $\mathcal{I}$ in the appropriate norm.

Remark. 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). As in the ring case, this may be done without loss of generality, since solving the $\mathcal{A}\text{-SVP}$ problem on the fractional ideal $\mathcal{I}$ may be done by solving it on the integral ideal $N_{K/\mathbb{Q}}(c)\mathcal{I}$, where $c \in \mathcal{O}_K$ is the element such that

$c\mathcal{I}$ is integral, and rescaling the solution.

Essentially, and unsurprisingly, we once again 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’ $a_1, \dots, a_N$ to be linearly independent over $\mathbb{Z}$. For BDD, we need a suitable ambient space and use the following definition.

Definition 3.5. Let $\mathcal{A}$ be a cyclic algebra, let $\mathcal{I}$ be some (possibly fractional) ideal of the natural order Λ , and let $\delta < \lambda_1(\mathcal{I})/2$. Then the $\mathcal{A}$ -BDD $_\delta$ problem, on input $y = x + e$ for $x \in \mathcal{I}$ and $e \in \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ satisfying $|e| \leq \delta$, is to compute x .

Although we restrict our discussion to the case where $|\sigma_i(\gamma)| = 1$ for each i , one could still define these problems for other choices of γ as long as a suitable norm is also defined. However, since this norm can no longer be defined using the Mod map we view our selection of norms as the correct ones for our needs.

3.2.1 Hardness of Order-Ideal Lattice Problems

The lattice problems we have chosen are not standard lattice problems, but as in the case of RLWE and Module LWE we assume that the problems are sufficiently hard for cryptographic purposes. Of course, analyzing this assumption is an important avenue of future research, but we can present some heuristic arguments as to why we suspect our assumption is likely to be correct. First of all, to date there are no efficient algorithms which solve short vector problems on ring ideal or module lattices despite substantial effort in this direction. A single u coordinate of an element of an ideal $\mathcal{I}$ does not in general fully determine the element, and it follows that the coordinates of an order-ideal lattice are not as correlated as the embedded coordinates of a ring ideal lattice. Similarly, a two-sided ideal of a natural order in general need not permit a representation as $\bigoplus_{i=0}^{d-1} u^i \mathcal{I}$ for an ideal $\mathcal{I}$ of $\mathcal{O}_L$, in which case the search for a short vector could be solved over $\mathcal{O}_L$ instead. Consequentially, we view the appropriate order-ideal lattice problems as likely harder than those on a ring lattice. Conversely, they are more structured

than a module lattice and we expect the hardness of the problems to line up accordingly, such that the module lattice problems are harder than ours.

Additionally, there has been recent movement towards quantum attacks on ideal lattice problems taking advantage of the underlying algebraic structure not present in the module case, such as those of [BS15], [BS16], [CDPR16] or [CDW17]. However, the majority of these attacks rely on two properties of the unit group of $\mathcal{O}_K$. Firstly, the bulk of these attacks make use of quantum algorithms to solve variants of the Abelian Hidden Subgroup Problem and hence implicitly rely on its commutivity, a property not satisfied by units of a natural order Λ . Secondly, they rely on the fact that $\mathcal{O}_K$ has an easy to compute and known unit group, whereas in non-commutative algebras computing the unit group of Λ even in low dimension can be very challenging; see for example [OLB09] for a discussion this topic. Hence, we view the order-ideal lattice problems as substantially safer from quantum attacks.

It is less clear how much harder, if at all, the module lattice problems are than the order-ideal variants. However, since the current technique for classic, ring, and module LWE problems is to treat them as unstructured instances of the LWE problem as in the CoreSVP method of NewHope [ADPS16], we deem it fair to consider the CLWE problem in the same way. Thus we consider our observation that the CLWE problem is likely harder than the RLWE problem as a sufficient argument for its security, at least given the state-of-the-art methods of attacking lattice problems.

3.3 Design Choices for CLWE

3.3.1 Relation to Module LWE

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

negacyclic matrix generated by the polynomial a and $\mathbf{s}, \mathbf{e}$ denote the coefficient vectors of s and e respectively. For an appropriate choice of error distribution this is precisely n LWE samples, with the exception that there is some structure in the matrix $\mathbf{A}$. By ordering the multiplication as $a \cdot s$ we get a similar transform from CLWE to MLWE. Assuming for now that we have a discretized form of CLWE and observing that for $q \in \mathbb{Z}$ we have $\Lambda_q \cong \bigoplus_{i=0}^{d-1} u^i \mathcal{O}_L / q \mathcal{O}_L$ (once again see [OAS12]), we transform a CLWE sample $a \cdot s + e$ into matrix-vector form to get $\phi(a) \cdot \mathbf{s} + \mathbf{e}$, where $\mathbf{s}$ and $\mathbf{e}$ are vectors of dimension d over $\mathcal{O}_L / q \mathcal{O}_L$. Setting $\mathbf{A} = \phi(a)$, one can see that for appropriate choices of error distribution this is equivalent to d samples from the MLWE distribution with the exception that there is now structure in the matrix $\mathbf{A}$, as intended. Of course, one could order the multiplication as $s \cdot a$ and we view it likely that the same (or similar) security properties could be found for a CLWE sample defined in this way, since we use the matrix embedding ϕ primarily as an intuitive tool rather than a technical one. On the other hand, some proofs would have to be substantially adapted and so we do not claim security of CLWE samples in this alternative ordering.

3.3.2 Natural Orders vs. Maximal

It is not immediately clear whether it is better to use the natural order or a maximal order in the definition of CLWE. Unlike in the case of RLWE we must make a choice; in a number field the ring of integers $\mathcal{O}_K$ is both the unique maximal order and 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 not trivial. Additionally, the natural order is suitably ‘orthogonal’ in that it has the same span in each u^i coordinate independent of the other coordinates. This is advantageous when considering the relation to module LWE, where the module taken is the full module $\mathcal{O}_K^d$, and also provides a significant advantage in terms of conceptual simplicity. However, there may be advantages to considering such constructions over a maximal order, perhaps by choosing specific number fields with well-behaved maximal orders. In particular, the two-sided ideal structure of an arbitrary maximal order is the same as that of the ring of integers of a number

field. To show search CLWE security in this case one would have to adapt the technical lemmas of the next section to the case of two-sided ideals of a maximal order, and due to their simpler structure these lemmas would likely be straightforward. However, we were unable to find a valid search-decision reduction for the case of maximal orders, since our methods rely on facts about extensions of ideals of the center to those of the natural order.

Another potential advantage of a maximal order in terms of concrete security relates to its decoding radius. Approximately speaking, the hardness of the LWE problem derives from viewing it as a decoding problem on the lattice generated by the appropriate order. This means that a denser lattice corresponds to a harder decoding problem and so a harder LWE problem. Since a maximal order is, by definition, maximally dense as a lattice it follows that intuitively the CLWE problem would be harder in a maximal order. To see this more rigorously note that any order, including the natural order, can be extended to a maximal order (see Corollary 10.4 of [Rei75]) on which we expect the decoding problem to be harder.

On the other hand, if one were to consider maximal orders one would also have to consider how to choose the maximal orders. There is no known efficient, general, algorithm for constructing a maximal order in a division algebra, or indeed a cyclic division algebra, in high dimension. Of course this need not necessarily be a problem, although it should be a consideration. The CLWE assumption over a maximal order would most likely use a fixed maximal order, so the object could be chosen amongst all maximal orders and pre-computed in a pre-processing phase, and not as a part of an encryption and decryption process. Alternatively, a family of algebras could be chosen with a maximal order that is easy to represent and construct. Overall, we would not consider this a large barrier to the use of maximal orders.

A final interesting property of choosing a maximal order is that one could drop the requirement that $\gamma \in \mathcal{O}_K$ and instead choose it to be any non-norm element of K . This would lead to additional flexibility in concrete constructions of cyclic division algebras, since enforcing γ to be integral and all its embeddings to have absolute value 1 has restricted us to extensions with non-norm elements which are also roots of unity. We see no way to make use of this observation, but it is plausible that dropping the restriction on γ may serve some practical purpose.

3.3.3 Security from Two Number Fields

In Module LWE, 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 a well-defined cyclic division algebra with appropriate γ . From a theoretical standpoint it is not immediately clear whether we want to consider asymptotic security in terms of d or n . Following our motivation from module LWE we suggest that n is likely the suitable choice, since the module rank d is typically small in applications using MLWE whereas the dimension of the underlying field K is large. However, there seems to be no a priori reason why with the right techniques one could not consider both n and d asymptotically; the only case our construction precludes is high dimensional Module LWE over a low dimension number field L , because the parameter d occurs in both the module and field dimension. However, our search-decision proof is only asymptotically efficient in the dimension of the base field K and so a new methodology would be required to show security for large d .

3.3.4 Requiring the Division Property

We have required that γ be chosen as a non-norm element so that $\mathcal{A}$ is a division algebra. We do this to avoid a ‘parallelizing’ attack on a protocol that was intended to provide a similar cyclic module LWE structure which we discuss now. For $m = 2$, the m -RLWE protocol of [PUTPPG16] can be considered as a structured variant of MLWE where the matrix $\mathbf{A}$ in the operation $\mathbf{A}\mathbf{s} + \mathbf{e}$ is a negacyclic matrix over some ring R_q . More explicitly, 2-RLWE considers the tensor product of two fields $K = K_1 \otimes K_2$ and runs the LWE assumption in the ring of integers R_q . The example use case given in [PUTPPG16] 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 $R_q = \mathbb{Z}_q[x, y]/[(x^{k_1} + 1), (y^{k_2} + 1)]$ 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 $\mathbb{Z}_q[x]/(x^{k_1} + 1)$ of module rank k_2 .

However, the observation of [BCV18] is that for $k_1 \geq k_2$ 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, [BCV18] 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 such a homomorphism removing the module structure of CLWE by taking a well-defined cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ to a smaller subfield containing K . Since any subfield is contained in at least one maximal subfield we restrict our search to maximal subfields of $\mathcal{A}$. It is a well known result on division algebras that any maximal subfield E of $\mathcal{A}$ contains K and satisfies $[E : K] = d$, and that in the case of a cyclic division algebra $\mathcal{A}$ there is a choice of $u' \in \mathcal{A}$ such that the cyclic algebra $\mathcal{A}' := \bigoplus_{i=0}^{d-1} u'^i E$ is isomorphic to $\mathcal{A}$ (see Section 15.1, Proposition A of [Pie82]). Assume for a contradiction that we had such a homomorphism $\chi : \mathcal{A} \rightarrow L$, 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 [BCV18] does not threaten our algebraic structure. Furthermore, a follow up paper [PUTPG⁺19] on m -RLWE observes that the attack of [BCV18] relies implicitly on the existence of zero divisors in the ring K , which are clearly excluded in the case where $\mathcal{A}$ is a division algebra.

Though this argument partially justifies the choice of non-norm γ , we have not shown the converse. That is, we have not presented a concrete attack on the CLWE problem in the case where the algebra is not a division algebra. We will return to this later after we have presented the primal form of CLWE, demonstrating some potential weaknesses of poorly instantiated CLWE.

3.3.5 CLWE as a One-Sided Problem

In the case of Ring-LWE, [LPR10] provides a reduction from short vector problems over ideal lattices to the problem of solving RLWE. However, since these reductions are only one-way, they do not say that the ability to solve ideal SVP implies an attack on the RLWE problem. Instead, the reduction implies only that the RLWE problem is *at least* as hard as ideal SVP. Similarly, we do not believe that an attack on the two-sided ideal lattice problems defined in Section 3.2 necessarily implies an attack on CLWE, and our reductions will show only that search CLWE is at least as hard as order-ideal SVP.

In fact, the gap between the security of the CLWE problem and ideal lattice problems over two-sided ideals of Λ may be substantial. This is because the CLWE problem is not naturally a two-sided problem as the underlying cyclic algebra is non-commutative. Instead, a single CLWE instance appears as a decoding problem on the right ideal generated by a , since $b = a \cdot s + e$ is close to precisely one element in the right span of a as long as e is small, $b \approx a \cdot s$. This is distinct from the ring case, where the commutativity means that the right and left spans of a align and are equal to $\langle a \rangle$. In our case, the two-sided ideal generated by a is a subset of the right ideal it generates. Alternatively, a collection of many samples from the CLWE distribution correspond to a ‘hidden’ left lattice problem, where one has to identify the hidden element s such that each b_i is close to the left span of s . Once again, this is naturally a one-sided problem rather than a two-sided one. In this way, we suspect that an effective attack on two-sided ideal problems need not lift to an effective attack on the CLWE problem.

The lattice problems of Section 3.2 are easy to generalize to the case of arbitrary ideals of Λ , and our restriction to two-sided ideals was artificial. However, this restriction was made to permit the security reductions which comprise the bulk of this section. The standard technique for transforming a BDD instance on an ideal lattice to a sample from an algebraic LWE distribution requires one to be able to ‘clear out’ ideal quotients, as in Lemma 3.7. The complexity of the required machinery to perform ideal quotients for one-sided ideals makes constructing such an operation substantially easier in the two-sided case, which is the reason we have restricted our lattice problems. However, we remark that lattice problems defined over right or left ideals

would be at least as hard as than those over two-sided ideals, because the set of two-sided ideals is a subset of the set of right ideals and the set of left ideals¹. We suggest constructing a security reduction for the CLWE problem basing the security on more general ideal problems as an interesting avenue of future research because it would likely more carefully quantify the security of the CLWE problem. Furthermore, the existence of one-sided ideals is one of the more substantial novel properties arising from the jump to non-commutative rings, and consequentially such a security reduction would be an object of theoretic interest.

3.4 Technical Lemmas

We now have our target lattice problems for CLWE security reductions and have discussed the choices made in the formalization of the CLWE problem. Before we can proceed with our reductions, we require some technical lemmas. Though the statements and proofs of these are novel, there are corresponding number field results shown in Section 2 of [LPR10] which are required in the analogous steps of the proof of RLWE security.

The lemmas will be required later to construct an efficiently invertible bijection between quotient spaces $\mathcal{I}/\langle q \rangle \cdot \mathcal{I}$ and $\Lambda/\langle q \rangle$, which is sometimes referred to as ‘clearing out’ the ideal $\mathcal{I}$. For the first lemma, one can prove a more general statement where $\mathcal{J}$ is not restricted to being generated by a prime integer, but more sophisticated algebraic machinery is required.

Lemma 3.6. *Let $\mathcal{I}$ be a two-sided ideal of the natural order Λ and let $\mathcal{J} = q \cdot \Lambda = \langle q \rangle \cdot \Lambda$, where q is a prime integer and $\langle q \rangle = \prod_{i=1}^r \mathfrak{q}_i^{e_i}$ is a decomposition into prime ideals in $\mathcal{O}_K$. Assume $\gamma \notin \mathfrak{q}_i$ for each i . Then, there exists an element $t \in \mathcal{I} \cap \mathcal{O}_K$ such that the ideal $t \cdot \mathcal{I}^{-1} \subset \Lambda$ is coprime to $\mathcal{J}$, and we can compute such a t efficiently given $\mathcal{I}$ and the prime factorization of $\mathcal{J}$.*

Remark. The condition on γ will be immaterial in our use case, since when γ is a unit the only $\mathcal{O}_K$ ideal that contains γ is $\mathcal{O}_K$ itself.

¹In fact, the set of two-sided ideals is the intersection of that of right and of left ideals.

Proof. For an ideal $\mathcal{I}$ denote by $\bar{\mathcal{I}}$ its intersection with K , which is a non-trivial ideal of $\mathcal{O}_K$ (see Section 3 of [OAS12]). We apply the corresponding Lemma 2.14 of [LPR10] to obtain $t \in \bar{\mathcal{I}}$ such that $t \cdot \bar{\mathcal{I}}^{-1}$ and $\bar{\mathcal{J}}$ are coprime as ideals of $\mathcal{O}_K$ and $t \in \bar{\mathcal{I}} \setminus \bigcup_{i=1}^r \mathfrak{q}_i \cdot \bar{\mathcal{I}}$. Assume, for a contradiction, that $t \cdot \mathcal{I}^{-1} + \mathcal{J} \neq \Lambda$ e.g. the ideals are not coprime. Then, there is some maximal two-sided ideal $\mathcal{M}$ of Λ containing $\mathcal{J}$ and $t^{-1} \cdot \Lambda$. By Propositions 1 and 4 of [OAS12], this ideal must be one of the ideals $\mathfrak{q}_i \cdot \Lambda$ since it contains $\mathcal{J}$. Then $t \cdot \mathcal{I}^{-1} \subset \mathfrak{q}_i \cdot \Lambda$ and consequentially $t \in \mathfrak{q}_i \cdot \bar{\mathcal{I}}$. Since t and $\mathfrak{q}_i$ are central it follows that $t \in \mathfrak{q}_i \cdot \bar{\mathcal{I}}$, a contradiction. $\square$

The next lemma will be the one required in our reduction. As in RLWE, in practice we are interested in the special case where $\mathcal{J} = \langle q \rangle$ for a prime integer q .

Lemma 3.7. *Let $\mathcal{I}, \mathcal{J}$ be two-sided ideals of Λ with $t \in \mathcal{I} \cap \mathcal{O}_K$ such that $t \cdot \mathcal{I}^{-1}$ and $\mathcal{J}$ are coprime ideals and let $\mathcal{P}$ denote an arbitrary fractional two-sided ideal of Λ . Then, the function $\chi_t : \mathcal{A} \rightarrow \mathcal{A}$ defined as $\chi_t(x) = t \cdot x$ induces a module isomorphism from $\mathcal{P} / \mathcal{J} \cdot \mathcal{P} \rightarrow \mathcal{I} \cdot \mathcal{P} / \mathcal{I} \cdot \mathcal{J} \cdot \mathcal{P}$. Furthermore, in the case where $\mathcal{J} = \langle q \rangle$ for a prime integer q we can efficiently (in probabilistic polynomial time) compute the inverse.*

Proof. The proof is similar to that of Lemma 2.15 of [LPR10]. Since t lies in the center of Λ it is clear that multiplication by t induces a module homomorphism. Given the map $\chi_t : \mathcal{P} \rightarrow \mathcal{I} \cdot \mathcal{P} / \mathcal{I} \cdot \mathcal{J} \cdot \mathcal{P}$ and $j \in \mathcal{J} \cdot \mathcal{P}$, we have $\chi_t(j) = t \cdot j \in \mathcal{I} \cdot \mathcal{J} \cdot \mathcal{P}$ and so it is clear that $\mathcal{J} \cdot \mathcal{P}$ is in the kernel of this map. Conversely, if $\chi_t(x) = 0$ then $t \cdot x \in \mathcal{I} \cdot \mathcal{J} \cdot \mathcal{P}$, from which it follows that $\mathcal{I}^{-1} \cdot t \cdot x \subset \mathcal{J} \cdot \mathcal{P}$. From the definition of coprime, $t \cdot \mathcal{I}^{-1} + \mathcal{J} = \Lambda$, and so there exists $a \in t \cdot \mathcal{I}^{-1}, b \in \mathcal{J}$ such that $a + b = 1$. Hence $x = (a + b) \cdot x = a \cdot x + b \cdot x$. Since $a \cdot x, b \cdot x \in \mathcal{J} \cdot \mathcal{P}$ it follows that $x \in \mathcal{J} \cdot \mathcal{P}$, from which injectivity follows immediately. As in the ring case, surjectivity follows from comparing norms.

To demonstrate efficient invertibility we must work slightly harder. Let $\mathcal{J} = \langle q \rangle$ and compute t as in the preceding lemma. The bijection $\chi_t : \Lambda_q \rightarrow \mathcal{I}_q$ is an additive homomorphism. Thus, it suffices to compute the inverse of all elements of a $\mathbb{Z}$ -basis of $\mathcal{I}_q$, since then any element can be inverted by computing its representation in this basis and inverting that. We construct such a basis as follows: first, choose $n^2 \cdot d^4$ elements $x_i, i = 1, \dots, n^2 \cdot d^4$ from Λ_q uniformly at random

and compute $y_i = \chi_t(x_i)$ for each i . It follows that each y_i is a uniformly random element of $\mathcal{I}_q$. With high probability the y_i 's form a spanning set of $\mathcal{I}_q$ (see Lemma 3.8), which we can reduce to a $\mathbb{Z}$ -basis $y'_1, \dots, y'_{n \cdot d^2}$ with the desired property that each element has a known inverse. If this algorithm fails (e.g. there is no suitable basis $y'_1, \dots, y'_{n \cdot d^2}$), we repeat, choosing a fresh set of elements $x_1, \dots, x_{n^2 \cdot d^4}$ until we succeed. $\square$

Lemma 3.8. *Given a set of $n^2 \cdot d^4$ independent and uniformly random elements $\Xi \subset \mathbb{Z}_q^{nd^2}$, the probability that Ξ contains no set of nd^2 linearly independent vectors (over $\mathbb{Z}$) is exponentially small in n .*

This lemma is a straightforward adaptation of Corollary 3.16 of [Reg05], but we present a proof for completeness.

Proof. Let $\Xi = \{X_1, \dots, X_{n^2 \cdot d^4}\}$. For $i = 1, 2, \dots, n \cdot d^2$ let B_i be the event

$$\dim \text{span}(X_1, X_2, \dots, X_{(i-1)n \cdot d^2}) = \dim \text{span}(X_1, \dots, X_{in \cdot d^2}) < n \cdot d^2.$$

If none of the events B_i occur then $\dim \text{span} \Xi = n \cdot d^2$, so it is sufficient to show that $\mathbb{P}[B_i]$ is exponentially small for each i . For fixed i , if

$$\dim \text{span}(X_1, \dots, X_{(i-1)n \cdot d^2}) \geq n \cdot d^2,$$

then B_i does not hold. Otherwise, $\dim \text{span}(X_1, \dots, X_{(i-1)n \cdot d^2}) < n \cdot d^2$ and so the samples $X_{(i-1)n \cdot d^2 + 1}, \dots, X_{in \cdot d^2}$ each lie in the span of $(X_1, \dots, X_{(i-1)n \cdot d^2})$ with probability at most $\frac{1}{q}$ (since we are working over $\mathbb{Z}_q$). Since they are chosen independently, the probability that none of them lie in the span is at most $\frac{1}{q^{n \cdot d^2}}$, which is exponentially small. $\square$

3.5 Hardness of Search CLWE

For the remainder of this section we will always work in a Galois extension of number fields L/K , where $[L : \mathbb{Q}] = [L : K] \cdot [K : \mathbb{Q}] = d \cdot n$, such that γ is a non-norm element and

$\text{Gal}(L/K) = \langle \theta \rangle$ is a cyclic group of order d . Recall from the motivation of structured MLWE that in practice we seek asymptotic security in n , since the parameter d corresponds to the typically small module dimension. Nonetheless, when considering the comparison to modules, our number fields have dimension nd . We abbreviate the condition $|\sigma_i(\gamma)| = 1$ for all i by $|\gamma| = 1$, and will do so for the rest of this work as well.

Definition 3.9. We define the family of error distributions Σ_α as the set of all Gaussian distributions D_Σ over $\bigoplus_{i=0}^{d-1} u^i L_\mathbb{R}$ with covariance matrix obtained as the distribution of the error in Lemma 3.16.

This is the family of error distributions that 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 each distribution in Σ_α is a Gaussian distribution whose marginals are Gaussian with variance at most α and whose covariance matrix has sparse off-diagonal terms.

In the following theorem we denote by $\mathcal{A}\text{-DGS}_\xi$ the problem of sampling a discrete Gaussian $D_{\mathcal{I},\xi}$, where $\mathcal{I}$ is some ideal of the natural order Λ . As usual, we denote by N the total dimension of $\mathcal{A}$, $N := nd^2$.

Theorem 3.10. *Let $\mathcal{A}$ be a cyclic division algebra over a number field L with center K and natural order Λ whose non-norm element satisfies $|\gamma| = 1$. Let $\alpha = \alpha(n) \in (0, 1)$ and $q = q(N) \geq 2$ be parameters such that $\alpha \cdot q \geq 2\sqrt{d} \cdot \omega(\sqrt{\log nd})$. Then, there is a polynomial-time quantum reduction from $\mathcal{A}\text{-DGS}_\xi$ to search $\text{CLWE}_{q,\Sigma_\alpha}$ for any $\xi = r \cdot \sqrt{d}\omega(\sqrt{\log(d \cdot n)})/\alpha q$, where $r > \sqrt{2}q \cdot \eta_\varepsilon(\mathcal{I})$.*

Here, $\text{CLWE}_{q,\Sigma_\alpha}$ denotes the search CLWE problem in the case where $s \in \Lambda_q^\vee$ and $\psi \in \Sigma_\alpha$ are not specified. From this we deduce the following corollary which is similar to the case of MLWE in [LS15], since the lattice structures of our algebra are merely special cases of their modules.

Corollary 3.11. *Let $\mathcal{A}, \Lambda, \alpha$ and q be as above. Then, there is a polynomial-time quantum reduction from $\mathcal{A}\text{-SIVP}_\xi$ to search $\text{CLWE}_{q,\Sigma_\alpha}$ for any $\xi = \sqrt{8Nd} \cdot (\omega(\sqrt{dn})/\alpha)$.*

The following theorem is our analog of the standard LWE bootstrapping step, such as Lemma 4.10 of [LS15].

Theorem 3.12. *Given an oracle that solves $\text{CLWE}_{q,\Sigma_\alpha}$ for 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 a sample from $D_{\mathcal{I},r'}$.*

We then prove Theorem 3.10 in the standard iterative manner for LWE problems as in [Reg05]. 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 theorem 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 $r'' \geq \xi$. In order to obtain the initial collection of Gaussian samples we use the sampler of Lemma 3.2 of [Reg05] on the module representation $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$. As usual, we obtain Theorem 3.12 in two steps, combining the main reduction of Lemma 3.16 with the following quantum step adapted from [Reg05]. We use a form of $\mathcal{A}\text{-BDD}_{L,\delta}$ from [LS15] where we bound the offset in the norm $\|e\|_{2,\infty} := \max_j \sqrt{(\sum_{i=0}^{d-1} |\sigma_j(e_i)|^2)} \leq \delta$, where σ denotes the canonical embedding of L .

Lemma 3.13. *There is an efficient quantum algorithm that given an $N = n \cdot d^2$ dimensional lattice $\mathcal{L} := \sigma_{\mathcal{A}}(\mathcal{I})$ corresponding to a two-sided ideal $\mathcal{I}$, a real $\delta < \lambda_1(\mathcal{L}^*)/(2\sqrt{2nd})$, and an oracle that solves $\mathcal{A}\text{-BDD}_{\mathcal{L}^*,\delta}$ with all but negligible probability, outputs an independent sample from $D_{\mathcal{L},\sqrt{d}\omega(\sqrt{\log(nd)})/\sqrt{2}\delta}$.*

3.5.1 BDD to Search CLWE

We proceed with the cyclic algebra analog of the BDD-to-LWE samples transformation from Section 4 of [LPR10]. As is standard for LWE security, we use the following ‘modulo q ’ definition of BDD:

Definition 3.14. For any $q \geq 2$, the $q\mathcal{A}\text{-BDD}_{\mathcal{I},d}$ problem is as follows: given an instance of the $\mathcal{A}\text{-BDD}_{\mathcal{I},\delta}$ problem $y = x + e$ with solution $x \in \mathcal{I}$ and error $e \in \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ satisfying

$$\|e\|_{2,\infty} \leq \delta, \text{ output } x \bmod q\mathcal{I}.$$

We use (a special case of) Lemma 3.5 from [Reg05], which lifts immediately since it is lattice preserving.

Lemma 3.15. *For any $q \geq 2$ there is a deterministic polynomial time reduction from $\mathcal{A}$ -BDD $_{\mathcal{I},d}$ to $q\mathcal{A}$ -BDD $_{\mathcal{I},d}$.*

That is, it is sufficient to solve the BDD instance modulo the integer q . 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 [LPR10], but has some technical differences induced by the structure of cyclic algebras.

Lemma 3.16. *Let $\mathcal{A}$ be as in Theorem 3.10. Then, there is a probabilistic polynomial time algorithm that on input a prime $q \geq 2$, order-ideal $\mathcal{I}^\vee \subset \Lambda$, a $q\mathcal{A}$ -BDD $_{\mathcal{I}, \alpha q \cdot \omega(\sqrt{\log(nd)})/\sqrt{2nd} \cdot r}$ 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 $D_{\mathcal{I},r}$ outputs samples that are within negligible statistical distance of the CLWE distribution $\Pi_{q,s,\Sigma}$ for a secret $s = \chi_t(x \bmod q\mathcal{I}^\vee) \in \Lambda_q^\vee$, where χ_t is as in Lemma 3.7 and Σ is an error distribution such that in the case where $|\gamma| = 1$ the resulting error e'' has marginal distribution in each coordinate that is Gaussian with parameter $r_{i,j} \leq \alpha$.*

Proof. The proof will be in two parts - first, we describe the algorithm, then we will demonstrate correctness. Recall that in the definition of the CLWE distribution a sample is in the form $(a, b) = (a, (a \cdot s)/q + e \bmod \Lambda^\vee)$, where e is taken from an error distribution $\psi \in \Sigma_\alpha$.

Begin by using Lemma 3.6 to compute an element $t \in \mathcal{I} \cap \mathcal{O}_K$ such that $\mathcal{I}^{-1} \cdot \langle t \rangle$ and $\langle q \rangle$ are coprime. 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 form the pair

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

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

We 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 $r \geq q \cdot \eta(\mathcal{I})$, so by appealing to Lemma 2.7 it can be seen that any value $z \bmod q\mathcal{I}$ 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 $z \bmod q\mathcal{I}$ and the uniform distribution is bounded above by 2ε . Since χ_t and its inverse are both bijections we conclude that $a = \chi_t^{-1}(z \bmod q\mathcal{I})$ is within statistical distance 2ε of the uniform distribution over Λ_q .

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

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

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

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

We can explicitly calculate the error as

$$e'' = \sum_{i=0}^{d-1} u^i \left(\sum_{j+k=i} \theta^k(z_j) \cdot e_k (1 - (1 - \gamma) \mathbb{1}_{j+k \geq d}) \right) + e' \quad (3.1)$$

where the sum $j+k$ is taken modulo d and the function $(1 - (1 - \gamma) \mathbb{1}_{j+k \geq d})$ is 1 if $j+k < d$ and γ otherwise². Since $|\gamma| = 1$ and $z \leftarrow D_{\mathcal{I},r}$ is spherically distributed it follows that multiplying by γ and applying the permutation on the σ_L coordinates of z_j 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 [LS15]. $\square$

²This term is just indicating whether or not we have had to use the relation $u^d = \gamma$ in this summand or not.

Though we do not specify the covariance of Σ , one can see that each entry of $\sigma_{\mathcal{A}}(z)$ appears in $\sigma_{\mathcal{A}}(e'')$ exactly d times. So, by symmetry, each element of $\sigma_{\mathcal{A}}(e'')$ has non-zero correlation with at most d^2 other entries. Consequentially, a proportion of at most $\frac{nd^4}{n^2d^4} = \frac{1}{n}$ of entries of Σ are non-zero. That is to say that although the coordinates are correlated, the correlations can be viewed as minor and sparse. In practice, we will ignore this correlation for convenience's sake and propose cryptosystems with error terms that are spherical Gaussians.

3.6 Search To Decision Reduction

In this section we will show that the hardness of decision CLWE follows from that of the search problem. However, as a function of our methods we will be forced to restrict the secret s in the search problem to some subset $G \subset \Lambda_q^\vee$ of size q^{nd} whose composition we discuss shortly. Once again, we follow a combination of the expositions of [LPR10] and [LS15] for the ring and module cases, making necessary changes for the structure of cyclic algebras. We will make heavy use of the CRT-style decomposition of [OAS12]. To keep our exposition explicit, and also to remind the reader, we state it in the form we use it in this section.

Lemma 3.17. *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 $\mathcal{I} = \mathfrak{q}_1 \dots \mathfrak{q}_n$ in $\mathcal{O}_K$. Then, we have the isomorphism*

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

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

In the case where γ is a unit, $\Lambda^\vee = \bigoplus_i u^i\mathcal{O}_L^\vee$. The $\subset$ inclusion follows by considering the trace of individual coordinates e.g. $x \cdot u^i\mathcal{O}_L$ for $x \in \Lambda^\vee$, and the $\supset$ inclusion holds since $\gamma\mathcal{O}_L = \mathcal{O}_L$ and the trace is additive. Hence, the above lemma is also valid in the case where each instance of $\mathcal{O}_L$ and Λ are replaced with their respective duals.

As in [LPR10], 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 certain families of algebras $\mathcal{A}$, such as those in Theorem 4.3, but since we have not formalized those algebras yet we state the exact conditions we need: that the modulus splits completely in the base field K , which is some cyclotomic field $K = \mathbb{Q}(\zeta_m)$ of dimension n , and that the non-norm element γ is an n^{th} root of unity. It follows that if $q \equiv 1 \pmod{m}$ then q splits completely into a product of prime ideals $\mathfrak{q}_1 \dots \mathfrak{q}_n$ as an ideal of $\mathcal{O}_K$, and we obtain the decomposition

$$\Lambda/q\Lambda \cong \mathcal{R}_1 \times \dots \times \mathcal{R}_n$$

where $\mathcal{R}_i$ is as in Lemma 3.17. In this section we will abuse notation and denote by $x \pmod{\mathcal{R}_i}$ the value of $x \in \Lambda^\vee/q\Lambda^\vee$ in the $\mathcal{R}_i$ coordinate under the isomorphism of Lemma 3.17.

Also as in [LPR10] we see no way to avoid randomizing the error distribution in the resulting decision problem. However, we also 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 variants of the CLWE problem. Namely, we assume that an oracle for a CLWE problem on an algebra $\mathcal{A} = (L/K, \theta, \gamma)$ is also an oracle for the problem on any algebra $\mathcal{A}' = (L/K, \theta, \gamma')$ over the same number fields L, K and some other primitive root of unity $\gamma' \in \mathcal{O}_K$. $\mathcal{A}'$ is also a cyclic algebra, since if K is cyclotomic then if some root of unity γ is a non-norm element then so are the other roots of unity γ' . We state this formally as follows.

Assumption. Let $\mathcal{O}$ be an oracle solving a CLWE problem in an algebra $\mathcal{A} = (L/K, \theta, \gamma)$ for some number fields L/K and primitive root $\gamma \in \mathcal{O}_K$. Then, $\mathcal{O}$ is also an oracle for the same CLWE problem in the algebra $\mathcal{A} = (L/K, \theta, \gamma')$ for the same number fields L/K and some other primitive root $\gamma' \in \mathcal{O}_K$.

Intuitively this assumption implies that for fixed L and K the hardness of the CLWE problem is invariant under the choice of root of unity γ , and will be required for Lemma 3.22. 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

isomorphisms sending $\mathcal{A}$ to $\mathcal{A}'$, which we will define shortly. Finally, we emphasize that our algorithm is only intended to be efficient in the dimension n of the base field K , and we expect to fix d as a small constant in practice.

3.6.1 Restricting the Secret Space

The decomposition of Lemma 3.17 will be an integral part of our reduction, since as in the ring and module case it allows us to ‘guess’ the value of the secret independently in each $\mathcal{R}_i$, which is of size polynomial in n . However, in Lemma 3.25 we will need to use a fact that is implicit in the search-decision reduction of [LPR10]: for uniformly random $v \in \mathcal{R}_i$ and an incorrect guess g of the secret s modulo $\mathcal{R}_i$, the distribution of $v(g - s)$ is uniformly random. In the ring and module cases, the secret space is decomposed into a direct product of finite fields, so it is clear that $v(g - s)$ is uniformly random in each finite field for $g \neq s$.

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

Definition 3.18. A set $G \subset \Lambda^\vee/q\Lambda^\vee$ is called a pairwise difference set if for $g \neq h \in G$ it holds for each i that $g - h$ is either invertible or 0 modulo each $\mathcal{R}_i$.

We also need to guarantee that there exist sufficiently large choices of G . A simple method for constructing a valid G_i is by fixing some arbitrary embedding β of $\mathbb{F}_{q^d}$ into $M_{n \times n}(\mathbb{F}_q)$ and letting G_i equal the image of this embedding, such that $|G_i| = q^d$ and $|G| = q^{nd}$. Indeed, a

G_i constructed in this way is maximal because any set of matrices in $M_{d \times d}(\mathbb{F}_q)$ of size at least $q^d + 1$ contains two matrices with the same first row, whose difference is therefore not invertible.

There are a number of choices of embedding β , and thus set G_i , equal to the number of irreducible polynomials of degree d in $\mathbb{F}_q[x]$. This can be calculated by the Necklace polynomial, and in general will vastly exceed q . We make clear that our reduction will take the decision CLWE problem for *arbitrary secret* s to the search CLWE problem where $s \in G$ for *arbitrary fixed* G , which we denote by $\text{CLWE}_{q, \Sigma_\alpha, G}$. Thus, our reduction states that the decision problem is as hard as the search problem for the hardest choice of G , precluding obvious attacks on the unique case where $G = \mathcal{O}_{Lq}^\vee$ and the CLWE problem with $s \in G$ corresponds to d parallel copies in L of the RLWE problem³. For a general set G , $s \in G$ will not provide parallelization since they need not have the property of L that they are entirely contained in one u coordinate of $\mathcal{A}$. Additionally, even though elements of G constructed this way co-commute, they do not lie in the center of Λ and the multiplication $a \cdot s$ in the CLWE instance will not be a commutative operation.

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

In fact, there is no a priori reason why G_i should be a field, or even closed under multiplication. For example, fixing a pair of invertible matrices M_1, M_2 and replacing G_i with $M_1 \cdot G_i \cdot M_2 = \{M_1 X M_2 \mid X \in G_i\}$ results in a new set of size q^d whose pairwise differences are all invertible but is not multiplicatively closed in general. Although the field embedding technique is perhaps the most elegant way of building G_i , and certainly the most constructive, it may transpire that taking s from some set with less algebraic structure is advantageous in terms of the hardness of the resulting search problem. One can also construct the valid set $G_i + X$ by adding a

³Although this case exists only when each $\mathfrak{q}_i \mathcal{O}_L$ is a prime ideal in $\mathcal{O}_L$.

fixed matrix X to each element of G_i , but this technique is somewhat constrained by the fact that LWE samples are additive in the secret s (e.g. one could just add $a \cdot X$ into the second coordinate of the resulting samples).

More generally, assuming further that elements of G_i are invertible, the set G_i corresponds to a rank metric code over $\mathbb{F}_q$ with maximum rank distance (MRD), a code in which the difference between any two elements has full rank [Gab85]. MRD codes seem to correspond quite naturally to cyclic division algebras, since they also permit a structured matrix representation that relates to invertible matrices. Classifying families of rank metric codes appears to be an open area of research, but [She16] contains a comprehensive summary and [LN16] discusses the automorphism groups of such codes, generalizing the notion of multiplying sets on the left and right by invertible matrices X and Y . We remark that this natural correspondence between rank metric codes and cyclic algebras may be an interesting avenue of future research, especially since the rank metric appears in a second round NIST process candidate, RQC [MAB⁺19].

Although this restriction is not ideal, we have two remarks about the implications on the security of the CLWE problem. Firstly, it is not uncommon for cryptosystems to have a security reduction to well known problems for the search variant, then argue the hardness of decision by observing that attacks on the search problem are the best known attacks on the decision version as well. For example, the second round NIST candidate RQC [MAB⁺19] is based on the hardness of the Ideal Rank Syndrome Decoding problem, for which the search problem can be reduced to the Hamming setting but the decision problem has no known hardness reduction. Secondly, restricting the secret space in (R)LWE problems is not an uncommon idea: tertiary secrets, where each coordinate of $s \in \{-1, 0, 1\}$, are used in the NIST candidate LAC [LLZ⁺18] amongst others, and security whilst restricting the secret to orders or subfields is discussed in [BBPS18], and to other K -lattices in [PP19]. Overall, we suspect that the decision CLWE problem is polynomial time equivalent to the search CLWE problem without restriction on s , in particular when the number of samples is small as in our applications in Chapter 5, and that the restriction is a function of our reduction technique rather than some causal property of the CLWE distribution. For the purposes of constructing a cryptosystem, we assume that this reduction implies that the decision CLWE problem is hard.

3.6.2 Outline of the Reduction

Definition 3.19. The error distribution Υ_α on the family of possible error distributions is sampled from by choosing an error distribution $\Sigma \leftarrow \Sigma_\alpha$, increasing the diagonal coordinates of Σ to α , and adding it to $D_{\mathbf{r}}$, where each $r_i := \alpha((n \cdot d^2)^{1/4} \cdot \sqrt{y_i})$ for $y_1, \dots, y_{n \cdot d^2}$ sampled from $\Gamma(2, 1)$.

Theorem 3.20. *Let Λ, q, L, K be as above with $q \in \text{poly}(n)$ and assume that $\alpha \cdot q \geq \eta_\varepsilon(\Lambda^\vee)$ for a negligible $\varepsilon = \varepsilon(n)$. Then, for any pairwise difference set $G \subset \Lambda_q^\vee$, there is a probabilistic reduction from $\text{CLWE}_{q, \Sigma_\alpha, G}$ to $D\text{-CLWE}_{q, \Upsilon_\alpha}$ which runs in time polynomial in n .*

We will prove Theorem 3.20 in the usual manner: first we show that it is sufficient to recover the value of $s \in \Lambda_q^\vee$ in one of the rings $\mathcal{R}_i$ (Lemma 3.22). Then, we use a hybrid distribution to define a decision problem in $\mathcal{R}_i$, for which we demonstrate a search to decision reduction (Lemma 3.25). We then use a hybrid argument to conclude the proof (Lemma 3.28).

3.6.3 CLWE in $\mathcal{R}_i$

Definition 3.21. The $\mathcal{R}_i\text{-CLWE}_{q, \Sigma_\alpha, G}$ problem is to find the value $s \bmod \mathcal{R}_i$ given access to the CLWE distribution $\Pi_{q, s, \Sigma}$ for some arbitrary $\Sigma \in \Sigma_\alpha$ and $s \in G'$ for some known, arbitrary pairwise difference set G' satisfying $|G| = |G'|$.

In the following lemmata we make use of the automorphisms of K coordinatewise on the rings $\mathcal{R}_i$. Since we assumed that 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 $\alpha_i : \mathcal{A} \rightarrow \mathcal{A}'$, with $\mathcal{A}' = (L/K, \theta, \gamma')$, which agree with α_i on L and send u to the new indexing element u' that satisfies relations $u'^d = \alpha_i(\gamma)$ and $xu' = u'\theta(x)$ for $x \in L$. By the assumption on K , $\alpha_i(\gamma)$ is a non-norm element since it is some d^{th} root of unity, and it follows 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 since each component of $z \cdot e + e'$ is defined coordinatewise over the d copies of $L_\mathbb{R}$ in the module representation of $\mathcal{A}$. Since α_i induces the same permutation of the entries of the canonical embedding of L in each coordinate it fixes the family of choices for each of z, e, e' ; now because α_i is homomorphic the family of distributions $z \cdot e + e'$ is fixed. It follows that the extended α_i function maps the $\mathcal{R}_i$ -CLWE $_{q, \Sigma_\alpha, G}$ 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 3.17) of $\Lambda_q^\vee$ by sending $\mathcal{R}_i$ to some $\mathcal{R}_j$, where j depends on the choice of σ_i , then applying some arbitrary isomorphism on $\mathcal{R}_j$. We are now ready for the first step of our reduction.

Lemma 3.22. *There is a deterministic polynomial time reduction from CLWE $_{q, \Sigma_\alpha, G}$ to $\mathcal{R}_i$ -CLWE $_{q, \Sigma_\alpha, G}$.*

Proof. Let $\mathcal{O}_i$ be an oracle for the $\mathcal{R}_i$ -CLWE $_{q, \Sigma, G}$ problem. Since Lemma 3.17 defines an isomorphism, it is sufficient to use $\mathcal{O}_i$ to solve the $\mathcal{R}_j$ -CLWE $_{q, \Sigma, G}$ problem for each j . Let $\alpha_{j/i}$ be an extension of the automorphism of K mapping $\mathbf{q}_j$ to $\mathbf{q}_i$, which exists by transitivity. Then, given a sample $(a, b) \leftarrow \Pi_{q, s, \Sigma_\alpha}$ with $s \in G$ we construct the new 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 $\mathcal{A}' = (L/K, \theta, \alpha_{j/i}(\gamma))$. By our assumption on CLWE oracles, feeding these samples into $\mathcal{O}_i$ outputs a value $t_j \bmod \mathcal{R}_i$.

We claim $\alpha_{j/i}^{-1}(t_j) = s \bmod \mathcal{R}_j$. Since $\alpha_{j/i}$ is an isomorphism, each sample (a, b) is mapped to a new CLWE sample $(\alpha_{j/i}(a), \alpha_{j/i}(a \cdot s/q + e) \bmod \Lambda^\vee)$ in a new algebra $\mathcal{A}'$, and we may write the second coordinate as $\alpha_{j/i}(a) \cdot \alpha_{j/i}(s)/q + \alpha_{j/i}(e) \bmod \Lambda^\vee$. Since our isomorphisms 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 $\alpha_{j/i}(s) \in \alpha_{j/i}(G)$ and error distribution $\Sigma' \in \Sigma_\alpha$. Since $\alpha_{j/i}$ permutes the rings $\mathcal{R}_i$ and is an isomorphism of $\Lambda_q^\vee$, $\alpha_{j/i}(G)$ is a valid pairwise difference set. Hence $\mathcal{O}_i$ outputs $t = \alpha_{j/i}(s) \bmod \mathcal{R}_i$, from which we recover $\alpha_{j/i}^{-1}(t) = s \bmod \mathcal{R}_j$, as required. $\square$

Remark. Here, our oracle finds the j coordinate in a new pairwise difference set G' obtained via an isomorphism, but it is enough for the next stage of our reduction that after applying $\alpha_{j/i}$

the secret still lies in some known, defined, pairwise difference set, even if the set has changed. That is, assuming that our solver works for an image $\alpha_{j/i}(G)$ in the $\mathcal{R}_i$ coordinate rather than G itself is not a concern.

3.6.4 Hybrid CLWE and Search-Decision

For this section we must introduce the cyclic algebra analog of the Hybrid LWE distribution used in [LPR10]; we use the decomposition into the rings $\mathcal{R}_i$ rather than the full Chinese Remainder Theorem of the ring case.

Definition 3.23. For a secret $s \in \Lambda_q^\vee$, distribution Σ over $\bigoplus_j u^j L_{\mathbb{R}}$, and $i \in [n]$, we define a sample from the distribution $\Pi_{q,s,\Sigma}^i$ over $\Lambda_q \times (\bigoplus_{j=0}^{d-1} u^j L_{\mathbb{R}})/\Lambda^\vee$ by taking $(a, b) \leftarrow \Pi_{q,s,\Sigma}$ and $h \in \Lambda_q^\vee$ which is uniformly random and independent mod $R_j, j \leq i$ and 0 mod $R_j, j > i$, and outputting $(a, b + h/q)$. If $i = 0$, we define $\Pi_{q,s,\Sigma}^0 = \Pi_{q,s,\Sigma}$.

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

Definition 3.24. For $i \in [n]$ and a family of distributions Σ_α , the W-D-CLWE $_{q,\Sigma_\alpha}^i$ problem is defined as the problem of finding j given access to $\Pi_{q,s,\Sigma}^j$ for $j \in \{i-1, i\}$ and valid CLWE secret and error distribution s, Σ .

Lemma 3.25. *For any $i \in [n]$ there is a probabilistic polynomial-time reduction from $\mathcal{R}_i$ -CLWE $_{q,s,\Sigma_\alpha,G}$ to W-D-CLWE $_{q,s,\Sigma_\alpha}^i$.*

Proof. We follow the standard search-decision methodology of guessing the value of the secret mod R_i and then modifying the samples so that the decision oracle tells us whether or not our guess was correct. Note that there are only $|G_i|$ possible values of $s \bmod \mathcal{R}_i$, which is bounded above by $|\mathcal{R}_i| = q^{d^2}$, and so we may efficiently (in n) enumerate over the possible values. Furthermore, the input to the $\mathcal{R}_i$ -CLWE $_{q,s,\Sigma_\alpha,G}$ problem defines the invertible difference set we must find s in.

We define the transform which takes a value $g \in G_i$ and maps $\Pi_{q,s,\Sigma}$ to $\Pi_{q,s,\Sigma}^{i-1}$ if $g = s \pmod{\mathcal{R}_i}$ or to $\Pi_{q,s,\Sigma}^i$ otherwise as follows. On input a CLWE sample $(a, b) \leftarrow \Pi_{q,s,\Sigma}$, output the pair

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

where $v \in \Lambda_q$ is uniformly random mod $\mathcal{R}_i$ and $0 \pmod{\mathcal{R}_j}$ for $j \neq i$ and $h \in \Lambda_q^\vee$ is uniformly random and independent mod $\mathcal{R}_j$ for $j < i$ and 0 on the other $\mathcal{R}_j$. 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 $g = s \pmod{\mathcal{R}_i}$, then $v(g - s) = 0 \pmod{\mathcal{R}_i}$ and so the distribution of the pair (a', b') is precisely $\Pi_{q,s,\Sigma}^{i-1}$. Otherwise, $v(g - s)$ is uniformly random mod $\mathcal{R}_i$ by assumption on G and $0 \pmod{\mathcal{R}_j}$, and so letting $h' = h + v(g - s)$ we see that the distribution of (a', b') is precisely $\Pi_{q,s,\Sigma}^i$. $\square$

Remark. This is the only stage of the proof which enforces that the asymptotic complexity scales only with n and not with d , since we are forced to guess all of $s \pmod{\mathcal{R}_i}$ at once. We do not see a way of separating the coordinates further, but observe that this restriction seems to be a function of our method and could plausibly be dropped by a superior technique. It is also the only stage of the proof which enforces that s lies in a pairwise different set. In fact, since the reduction is secret preserving the required decision oracle for $\text{W-D-CLWE}_{q,s,\Sigma_\alpha}^i$ could have the additional restriction that $s \in G$, but for the purposes of the rest of our proof it will be more convenient to have access to an oracle solving the more general problem where s is arbitrary. Additionally, in practical applications we will use the decision problem for arbitrary s , so we see no benefit of the tighter reduction where s is restricted.

3.6.5 Worst-Case to Average-Case Decision Reduction

Now that we have removed the restriction that $s \in G$, we are able to follow the skeleton of the RLWE search-decision reduction of [LPR10] more liberally.

Definition 3.26. For $i \in [n]$ and a distribution Υ_α over possible error distributions, an algorithm solves the $D\text{-CLWE}_{q, \Upsilon_\alpha}^i$ problem if with non-negligible probability over the choice of pairs $(s, \Sigma) \leftarrow U(\Lambda_q^\vee) \times \Upsilon_\alpha$ it has a non-negligible difference in acceptance probability on inputs from $\Pi_{q, s, \Sigma}^i$ and $\Pi_{q, s, \Sigma}^{i-1}$.

This is the average case decision problem relative to $\mathcal{R}_i$; 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 3.27. *For any $\alpha > 0$ and $i \in [n]$ there is a randomized polynomial-time reduction from $W\text{-}D\text{-CLWE}_{q, \Sigma_\alpha}^i$ to $D\text{-CLWE}_{q, \Upsilon_\alpha}^i$.*

Proof. Since the definition (Definition 3.19) 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 [LPR10], except that we replace each instance of $\text{mod } \mathfrak{q}_i R^\vee$ with $\text{mod } \mathcal{R}_i$ and each instance of R_q with Λ_q . $\square$

Remark. This choice of Υ_α means that the error in our decision problem is closer to diagonal than the corresponding error in the 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. This further justifies our desire to disregard the correlation of error coordinates in practical applications.

Finally, we conclude with our hybrid argument. We must first show that $\Pi_{q, s, \Sigma}^n$ is uniformly random given Σ sampled from Υ_α . Again this follows the same method as the ring case, except that we must appeal to the more general Theorem 3.1 of [Pei10] rather than Lemma 2.7.

Lemma 3.28. *Let Υ_α be as above and let $s \in \Lambda_q^\vee$. Then given an oracle $\mathcal{O}$ which solves the D - $CLWE_{q,\Upsilon_\alpha}$ problem there exists an efficient algorithm that solves D - $CLWE_{q,\Upsilon_\alpha}^i$ for some $i \in [n]$ using $\mathcal{O}$.*

Proof. The proof is identical to the ring case, Lemma 5.14 of [LPR10], except that the indexing set $\mathbb{Z}_m^*$ is replaced by $[n]$. $\square$

To conclude this section, we summarize our results by stating that we have shown that the hardness of Search CLWE follows from the presumed hardness of lattice problems derived from the embedding of two-sided ideals of the natural order Λ , and the hardness of the Decision CLWE problem for arbitrary secret follows from a restricted variant of the search problem. Although the location of the secret is heavily restricted in this search problem, we still assume that it is hard enough to imply that the decision problem is also hard. However, this hardness was only shown for the dual form and not the more practical discrete form. In Chapter 5 we will show how to mold the CLWE assumption into a form more suitable for cryptography by providing a primal and a normal form reduction. Before finishing this chapter, we follow a brief diversion and discuss an alternative method for demonstrating security of the decision RLWE problem, and why we were unable to adapt it to the case of CLWE.

3.6.6 Non-Applicability of the Direct to Decision Technique

Although the search reduction is agnostic about the choice of fields, the reductions demonstrating hardness of the decision problem for RLWE in [LPR10] only apply to a certain family of number fields: the search-decision step requires that the modulus q splits completely into a product of distinct prime factors in $\mathcal{O}_K$. This result was improved on in [PRSD17], where they present a direct reduction from the ideal lattice BDD problem to the decision RLWE problem for any ring and modulus. Following the usual reductions from ideal SVP problems to BDD on ideal lattices concludes a suitable argument for hardness of D-RLWE that holds for any K and q . Our search-decision reduction for CLWE comes with similar constraints to that of the ring case, since we have the same conditions on our center K and modulus q as they do on

their number field and modulus. Additionally, it is only valid for a restricted variant of the search problem. Consequentially, we attempted to adapt the techniques of [PRSD17] to the case of cyclic algebras. Unfortunately, we were unable to do so due to the structure of the multiplication operation in cyclic algebras. In this section we will briefly explain the techniques of [PRSD17] and why they do not easily lift to the CLWE case.

Essentially, the method of [PRSD17] is to perform a hill climbing algorithm on each $K_{\mathbb{R}}$ coordinate of the BDD offset e in $y = x + e$, using the oracle for the decision RLWE problem to guide the hill climbing. Specifically, they provide an algorithm to solve the following ‘Oracle Hidden Centre Problem’ (OHCP).

Definition 3.29. For parameters $\varepsilon, \delta \in [0, 1)$, the (ε, δ) -OHCP is an approximate search problem defined as follows. Given a scale parameter $d > 0$ and access to a (randomized) oracle $\mathcal{O} : \mathbb{R}^n \times \mathbb{R}_{\geq 0} \rightarrow \{0, 1\}$ such that its acceptance probability $p(\mathbf{z}, t)$ depends only on $\exp(t) \|\mathbf{z} - \mathbf{z}^*\|$ for some unknown ‘hidden center’ $\mathbf{z}^* \in \mathbb{R}^n$ with $\delta d \leq \|\mathbf{z}^*\| \leq d$, the goal is to output $\mathbf{z} \in \mathbb{R}^n$ satisfying $\|\mathbf{z} - \mathbf{z}^*\| \leq \varepsilon d$.

Remark. This definition extends easily into complex spaces by taking the natural embedding of $\mathbb{C}$ into $\mathbb{R}^2$.

To leverage the OHCP solver on a BDD instance they present the following specialized variant of the BDD instance to RLWE sample transformation, the ring equivalent of our Lemma 3.16.

Lemma 3.30. *There is an efficient algorithm that on input an integer $q \geq 2$ with known factorization, a fractional K -ideal $\mathcal{I}^\vee$, a coset $e + \mathcal{I}^\vee$, a bound $d \geq \|e\|_\infty = \max_i |\sigma_i(e)|$, a parameter $r \geq \sqrt{2}q \cdot \eta(\mathcal{I})$, and samples from $D_{\mathcal{I}, \mathbf{r}}$ for some $\mathbf{r} \geq r$, outputs samples that are within negligible statistical distance of the Ring-LWE distribution $A_{s, \mathbf{r}'}$ for a uniformly random $s \in \mathcal{O}_{K_q}^\vee$, where the coordinates of $\mathbf{r}'$ are given by $r'_i := (r_i |\sigma_i(e)|/q)^2 + (rd/q)^2$.*

The precise analysis of the coordinates of $\mathbf{r}'$ in the error distribution is both an integral part of their method and the reason why their techniques do not extend to cyclic algebras. The origin of this term comes from the error term e'' in the RLWE sample generated by the BDD step,

which is precisely

$$e'' = (z \cdot e)/q + e'$$

where $z \leftarrow D_{\mathcal{L}, \mathbf{r}}$ for a vector $\mathbf{r}$ with each $r_i \geq r$, e the offset in the BDD problem, and e' some extra additive error from a continuous Gaussian with variance rd/q . The crucial observation of [PRSD17] is that if one subtracts some guess $\bar{e}$ for e in the BDD instance before inputting it into the BDD transformation, then after the transform is applied the variance coordinate r'_i becomes

$$(r_i |\sigma_i(e - \bar{e})|/q)^2 + (rd/q)^2.$$

The next observation is that the success probability of a decision RLWE oracle, necessarily a probabilistic object, is not monotonically decreasing in the overall distance of the guess from true, $|e - \bar{e}|$. This is because the precise behaviour of a decision oracle for the RLWE problem is not necessarily symmetric about the coordinates and need not behave in an ‘intuitive’ manner, which is also true of arbitrary decision oracles for the CLWE problem as well. Thus, the OHCP algorithm cares about the distribution input to the oracle rather than the methodology of the oracle, since in this way as long as we only care about the center then controlling the distribution is enough.

The OHCP solver can be used to find the error e , and hence solve the BDD instance $y = x + e$, as follows. Solve the OHCP problem in 1 dimension for each coordinate $\sigma_i(e)$ separately by inputting $y + \bar{e}_i^*$ as a guess for $\sigma_i(e)$, where $\bar{e}_i^*$ satisfies $\sigma_j(\bar{e}_i^*) = \bar{e}^*$ if $i = j$ and 0 otherwise. Placing this guess into Lemma 3.30 gives samples whose error depends on $r_i |\sigma_i(e) - \bar{e}^*|$, and indeed only on this term since the other coordinates of the canonical embedding have not been altered. When the guess gets too close to being correct the oracle may lose sensitivity, so it is necessary to input the $\exp(t)$ term from Definition 3.29 into r_i to ‘scale up’ the error and recover sensitivity, but the other r_j all remain equal to r . A similar method is applied for the LWE case in [PRSD17] and can be applied to the MLWE case as well.

Unfortunately, the nature of the OHCP renders this method unsuitable for CLWE. The issue arises from the requirement that the OHCP has two inputs: a scale factor $\exp(t)$ and a guess of the center $\mathbf{z}^*$. Inputting the guess for the center, corresponding to the BDD offset, is not a concern: take the ideal-BDD instance $y = x + e$ and subtract some guess for one entry of the canonical embedding of some u^j coordinate of e without scaling. Upon inputting $y = x + e - \bar{e}$ into the CLWE transform, the resulting distribution depends only on the distance $|\sigma_i(e_j) - \bar{e}^*|$ as long as each coordinate $r_{i,j}$, where $z \leftarrow D_{\mathcal{I}, \mathbf{r}}$, is the same. Although how this distribution depends on this term is more complex than in the ring case, it still satisfies the necessary condition that it is dependent on only this term. Unfortunately, when inputting the scaling parameter $\exp(t)$ into $r_{i,j}$ we run into some issues. Unlike in the ring case, we do not have a convenient, coordinatewise, multiplication operation, which means that when we change the parameter $r_{i,j}$ from r to $r \exp(t)$ we do not get a distribution that depends only on $\exp(t)|\sigma_i(e_j) - \bar{e}^*|$. Although the resulting distribution does depend on this term, there is additional ‘noise’ input to the oracle by the effect of changing $r_{i,j}$ on other coordinates of the input error $(z \cdot e)/q + e'$, which occurs because the $\exp(t)$ term hits other coordinates in the multiplication step. Thus we do not have a valid instance of the OHCP to solve.

To see the problem more precisely, imagine a simplified case with a cyclic algebra $\mathcal{A} = (L/K, \gamma, \theta)$ with $K = \mathbb{Q}$ and $[L : K] = 2$, in which an element of $\mathcal{A}$ can be written $x_0 + ux_1$. Let $\mathcal{I}$ be an ideal of Λ and let $y = x + e$ be our BDD instance. The canonical embedding of L can be written (σ_1, σ_2) and so to define $D_{\mathcal{I}, \mathbf{r}}$ the vector $\mathbf{r}$ must have dimension 4. Say we wanted to guess the σ_1 coordinate of the u coordinate of e . Input a guess for this coordinate of the error into the BDD instance as $u\bar{e}_1$, where $\bar{e}_1 \in L_{\mathbb{R}}$ has $\sigma_1(\bar{e}_1) = \bar{e}^*$, our guess for $\sigma_1(e_1)$, and $\sigma_2(\bar{e}_1) = 0$. Then, inputting $y = x + e + u\bar{e}_1$ into the BDD to CLWE sample transformation results in an error term whose distribution depends only on the distance $|\sigma_1(e_1) - \bar{e}^*|$, the distance of our guess from the true value. However, since the oracle is not sensitive enough once we get close to the true result, we must input the scaling factor into the appropriate coordinate of $\mathbf{r}$. Now since

$$z \cdot (e - \bar{e}^*) = z_0(e - \bar{e}^*)_0 + \gamma\theta(z_1)(e - \bar{e}^*)_1 + u(z_1(e - \bar{e}^*)_0 + \theta(z_0)(e - \bar{e}^*)_1),$$

one can see that changing any coordinate of $\mathbf{r}$ causes an accidental ‘noisy’ scaling of a coordinate of e . For example, changing the obvious choice of the $r_{1,1}$ coordinate scales up the appropriate σ_1 coordinate of e_1 that we are trying to guess, but our distribution no longer depends only on the scaling of this guess as we also increased the σ_2 coordinate of e_1 in the u^0 coordinate of $z \cdot (e - \bar{e}^*)$ (since θ permutes σ_1 and σ_2). In this way, the methodology falls apart.

One could try to alter the use of the OHCP solver by guessing some subset of the coordinates of e at a time. Unfortunately, we were unable to separate e into subsets that are both disjoint under the multiplication operation and satisfy that CLWE oracles are monotonic in the distance of a guess from correct in this setting. Similarly, one could attempt to guess all of e at once, but the structure of CLWE multiplication means that the resulting distribution does not depend only on $\exp(t)|e - \bar{e}|$.

In order to adapt this method, one would need to find separable subsets of coordinates under the multiplication operation or equivalently find an orthogonal basis of $\mathcal{A}$. Otherwise, we are unable to use this method to show hardness of the decision CLWE problem for any cyclic division algebra and modulus.

Chapter 4

Concrete Algebras for CLWE

“That’s right,” shouted Vroomfondel,
“we demand rigidly defined areas of
doubt and uncertainty!”

The Hitchhiker’s Guide to the Galaxy

By Douglas Adams

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

- The non-norm element γ must lie in $\mathcal{O}_K$ to keep the natural order closed under multiplication, and should satisfy $|\gamma| = 1$ in order to maintain both the coordinatewise independence and sub-multiplicative properties of the norm¹.
- The dimension $n := [K : \mathbb{Q}]$ of the division algebra should be large and the degree $d := [L : K]$ should be small. This is to maintain the analogy with structured MLWE (the

¹Recall that we are abusing notation and writing $|\gamma| = 1$ for the condition $|\sigma_i(\gamma)| = 1$ as complex numbers for each L -embedding σ_i

degree corresponds to the module rank) and follows from the search-decision reduction, which takes time polynomial in n but not in d .

- The base field K should be cyclotomic and q should split completely in K . This is also a result of the methodology of the search-decision reduction, which uses the well understood factorization of $\langle q \rangle$ in $\mathcal{O}_K$. In addition, since the bulk of lattice based cryptography is done over cyclotomic fields, we consider algebras which are small extensions of these as somewhat natural. We observe that an improved proof of decision security may allow this point to be dropped, whereas the other two points feel more integral.

Although significant effort has been expended by coding theorists to construct cyclic division algebras satisfying a variety of conditions, such as in [VHLR09] or [LMM08], we find ourselves with a fairly unique set of restrictions. In particular, for reasons relating to desired applications, the majority of algebras used in coding theory are either of small total dimension or have small $[K : \mathbb{Q}]$ and scale asymptotically in $[L : K]$. Since we are interested in scaling up K asymptotically, we will have to build novel algebras satisfying the above requirements ourselves. We will, however, make heavy use of the following theorem from [LMM08] as an intermediate step. Here ζ_n denotes a primitive n^{th} root of unity.

Theorem 4.1. *Let $n = p^a$ be a prime power and let $K = \mathbb{Q}(\zeta_n)$. Then, there exist infinitely many cyclic Galois extensions L/K of degree n satisfying the condition that ζ_n^i is not a norm of L/K for $0 < i < n$.*

The proof of this theorem relies on the theory of localization of rings at prime ideals, which is beyond the scope of this thesis, and so we will not present it here. However, we remark that the theorem is effective in the sense that it provides an explicit description of L , and we provide a summary of the recipe for constructing L . The crucial aspect of its construction is that L is a subfield of some cyclotomic extension of K , $K(\zeta_s)$ for a prime s , but we present its full description for completeness.

First, find some prime² s such that $s \equiv 1 \pmod{p^a}$ but $s \not\equiv 1 \pmod{p^{a+1}}$, so that p^a is the

²It is easy to show that infinitely many primes satisfying this condition always exist by appealing to classical theorems of Chebotarev or Dirichlet.

highest power of p dividing $s - 1$. Set $M = K(\zeta_s)$ so that by coprimality $M = \mathbb{Q}(\zeta_{sn})$. Then $\text{Gal}(M/K)$ is a cyclic group of order $s - 1$ generated by some automorphism σ . Denote by L the subfield of M fixed by σ^n . Then $[L : K] = n$ by the fundamental theorem of Galois theory and the extension is both cyclic and Galois. Finally, localization theory is used to show that the powers of ζ_n are not norms in this extension. In this way, the theorem constructs L explicitly.

The part of this theorem of our interest is that it allows us to scale K asymptotically, but this comes with a drawback of high degree of L . We present a variety of methods that use this theorem as a starting point to construct good algebras satisfying our restrictions.

4.1 Impossible Algebras

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

Theorem 4.2. *Let $K = \mathbb{Q}(\zeta_m)$ for some positive integer m and let $p \geq 2$ be some integer which is coprime with m . Then, for any Galois extension L/K of degree p each $\zeta_m, \zeta_m^2, \dots, \zeta_m^{m-1}$ lies in $N_{L/K}(K^\times)$.*

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

This theorem precludes the existence of a very large class of cyclic division algebras with cyclotomic base field. In particular, if the degree of $[L : K]$ is coprime with m then we can not have our restrictions that $|\gamma| = 1$ and is integral whilst K is cyclotomic. We draw attention to the specific classes whose non-existence we are interested in: in an ideal world we might instantiate CLWE with $K = \mathbb{Q}(\zeta_{2^k})$ and $[L : K] = d$ for arbitrary small integer d corresponding to the module rank, which in practice is likely to be at most say 5. However, as a result of Theorem 4.2 we know that d can not be coprime with 2^k and must be even in order to permit a suitable γ , from which it follows that we can not have $d = 3, 5$.

In the next section we will find positive examples of explicit algebras for the cases which our result leaves open, such as $d = 2$ and $d = 4$, as well as valid extensions of odd module rank.

4.2 Suitable Algebras for Cryptography

Now that we know what d our search can be restricted to, we are ready to present a variety of methods to construct families of algebras that we suggest as suitable for cryptographic use. Our constructions will begin with Theorem 4.1 and then use elementary methods from Galois theory to build more favourable fields.

4.2.1 Valid Extensions of Prime Power Cyclotomics

We start with what we view as our best constructions, where we squash the field L from Theorem 4.1 to a subfield M of small index over the base K satisfying the necessary properties to generate a cyclic algebra.

Theorem 4.3. *Let $K = \mathbb{Q}(\zeta_n)$ be a prime power cyclotomic with $n = p^a$ for some integer a and prime p . Then, there exists a cyclic Galois extension M/K of any index d dividing n within which ζ_n satisfies the non-norm condition.*

Remark. Since the proof will provide an explicit description of M , the correct interpretation of this theorem is that we can construct cyclic division algebras $\mathcal{A} = (M/K, \theta, \gamma)$ with $\langle \theta \rangle =$

$\text{Gal}(M/K), \gamma = \zeta_n, K = \mathbb{Q}(\zeta_n)$, and $[M : K]$ is any divisor of $n = p^a$.

Proof. Let $K = \mathbb{Q}(\zeta_n)$ for a fixed $n = p^a$ with prime p and integer a . Following the construction of Theorem 4.1 fix a cyclic Galois extension L/K of degree n such that ζ_n^i is not a norm of an element of L into K for any $i = 1, 2, \dots, n-1$. We will choose M as a suitable intermediate extension $L/M/K$. Let σ denote the generator of $\text{Gal}(L/K)$, an automorphism of degree n . For d dividing n , σ^d fixes an extension M of K with $[L : M] = |\text{Gal}(L/M)| = n/d$ and it follows from the tower lemma that $[M : K] = d$. We will show that M is a satisfactory extension of K .

First, since $\text{Gal}(L/M)$ is a normal subgroup of $\text{Gal}(L/K)$ we see that M/K is a normal, and hence Galois³, extension. It follows from standard Galois Theory that

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


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

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

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

where the first equality follows from $m \in M$ and the second since the norm is multiplicative. $\bar{L}$ does not contain any power of ζ_n except $\zeta_n^n = 1$ since ζ_n is a non-norm element in L/K , so it follows that $n \nmid (n/d)i$ and so $d \nmid i$. From this we conclude that $\zeta_n, \zeta_n^2, \dots, \zeta_n^{d-1}$ do not lie in $\bar{M}$ and so ζ_n satisfies the non-norm condition. $\square$

³Since in this case all extensions are separable.

Figure 4.1: Cyclic Extensions Between L and K .

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

This is an effective construction that allows us to build cyclic division algebras of the form $\mathcal{A} = (M/K, \theta, \gamma)$ where $|\gamma| = 1$, K is an arbitrary prime power cyclotomic, and M is an extension of K with degree divisible by the prime p . For cryptographically relevant examples, we can consider degree 2 or 4 extensions of a 2-power cyclotomic or degree 3 extensions of a 3-power cyclotomic. Given the impossibility result of the previous section and the restriction on the absolute value of γ we view these algebras as essentially the best possible, at least for the case where K is a prime-power cyclotomic. Now that we have an appropriate method for constructing fields for CLWE we are almost ready to present some sample parameters. Before that, we briefly present another method of constructing algebras that allows us to target dimensions not achievable in this setting.

4.2.2 Constructions Using Compositums of Fields

Our other general method for constructing suitable extensions starts from extensions which are nearly what we are looking for and applies field compositums. We remark that building

cyclic algebras using compositums is a standard technique, and may be found for example in the section on cyclic algebras in Chapter 30 of [Rei75]. We recommend this method to build on top of fields constructed using either Theorem 4.1 or Theorem 4.3. We recall the definition of a field compositum.

Definition 4.4. Given a pair of number fields L and K , the field compositum LK denotes the smallest field containing both L and K .

It is easy to show that the compositum of two number fields exists and is well-defined by e.g. computing it as a subfield of a field E containing L and K . Say we have a Galois field extension E/M with non-norm element $\gamma \in \mathcal{O}_M$ whose Galois group is cyclic of degree d . Let F be some other Galois number field with $F \cap E = \mathbb{Q}$. Then $\text{Gal}(EF/MF) \cong \text{Gal}(E/M)$ and γ is a non-norm element in EF/MF . Relabelling this extension as L/K and letting θ denote the cyclic generator of the Galois group gives a cyclic field extension with non-norm γ such that $[L : K] = d$ and $[K : \mathbb{Q}] = [M : \mathbb{Q}] \cdot [F : \mathbb{Q}]$.

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

4.3 Sample Parameters

Now that we have discussed our techniques for constructing suitable number fields we proceed to demonstrate that these methods are able to attain cryptographically relevant dimensions. In this section, we present a small selection of proof-of-concept dimensions where we take our motivation for choices of dimension from Kyber and NewHope, since they are the successful first round NIST candidates whose methods are most similar to our own. Thus we aim for dimensions in the region of between 512 and 1024, dimensions proposed for both NewHope

and Kyber (which also achieves dimension 768). Of course, these schemes are restricted to having power-of-two ring dimension n and so their choices of dimension may not be optimal in general, but FrodoKEM [BCD⁺16], a plain LWE scheme, suggests dimensions in around the same range, specifically 640, 976, and 1344, so we consider dimensions in this region a sensible starting point. Corresponding to Kyber and other MLWE based schemes we will set a small ‘module’ rank $d := [\mathcal{A} : L]$. We are constricted in our choice of fields by the fact that d appears as a square in the total dimension $N = nd^2$, but for the most part we are able to work around this problem.

Method	Center	$[L : K]$	Total Dimension of $\mathcal{A}$
Theorem 4.3	$\mathbb{Q}(\zeta_{81})$	3	486
Theorem 4.3	$\mathbb{Q}(\zeta_{256})$	2	512
Theorem 4.3	$\mathbb{Q}(\zeta_{64})$	4	512
Theorem 4.3	$\mathbb{Q}(\zeta_{512})$	2	1024
Theorem 4.3	$\mathbb{Q}(\zeta_{128})$	4	1024
Theorem 4.3	$\mathbb{Q}(\zeta_{243})$	3	1458
Compositum	$\mathbb{Q}(\zeta_{192})$	3	576
Compositum	$\mathbb{Q}(\zeta_{576})$	2	768
Compositum	$\mathbb{Q}(\zeta_{384})$	3	1152

Table 4.1: Algebras Constructed in Section 4.3.

4.3.1 Two-Power Cyclotomic K

We begin with straightforward cases where we can apply Theorem 4.3 immediately to obtain fields in suitable dimensions. Let K be a two-power cyclotomic field, $K = \mathbb{Q}(\zeta_{2^k})$, with dimension $n := 2^{k-1}$. Since the rank $d = [L : K] = [\mathcal{A} : L]$ is a small power of two, the dimension n of K will be dictated by the choice of module rank d . We construct rank 2 and 4 examples as follows:

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

To obtain algebras in dimension 512 simply pick K with dimension $n/2$ e.g. $\mathbb{Q}(\zeta_{256})$ and $\mathbb{Q}(\zeta_{64})$ respectively. In all cases, Theorem 4.3 lets us pick the non-norm element γ as a root of unity.

4.3.2 Three-Power Cyclotomic K

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

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

4.3.3 Fields Using Compositum Techniques

The algebras with prime-power cyclotomic centers of the previous section use the field construction technique of Theorem 4.3, and as such they are restricted to algebras whose dimension N is in the form $p^k(p-1)$ for a prime p and integer k . In this section, we present examples of algebras with less restrictive dimension using field compositum techniques. We propose two alternate methods of applying field compositums: either use Theorem 4.3 to make an algebra which already has large dimension by selecting large center K and small extension L , then compose a small field onto K and L to tweak the total dimension. Alternatively, one can create algebras by selecting small fields L and K using Theorem 4.1 and composing both with a large field M .

We begin with an example of the first method that achieves dimension 768. Let L' be a degree two extension of the field $K' = \mathbb{Q}(\zeta_{64})$ chosen by Theorem 4.3 with non-norm root of unity γ , so that the corresponding algebra $\mathcal{A}'$ has dimension 128. Compose both L' and K' with the field $M = \mathbb{Q}(\zeta_9)$, denoting the compositums by L and K respectively. Then γ is still a non-norm element in the extension L/K , a degree two extension that is cyclic and Galois, and the algebra $\mathcal{A} = (L/K, \theta, \gamma)$ is a cyclic algebra of dimension $6 \times 128 = 768$, as required. We observe that here the center K corresponds to the fields with fast operations used in [LS19].

Our final method of composing large degree fields onto small degree extensions is aimed at targeting odd module ranks. Begin by choosing the desired module rank d as a (likely small) odd prime. Then set $K' = \mathbb{Q}(\zeta_d)$ and pick L' as a cyclic Galois extension of K' in which the d^{th} root of unity is a non-norm element using Theorem 4.1. Let $M := \mathbb{Q}(\zeta_{2^k})$ and again let L and K denote its compositum with L' and K' respectively. Then $\mathcal{A} = (L/K, \theta, \gamma)$ is a cyclic algebra with $n := [K : \mathbb{Q}] = (d-1)2^{k-1}$ and $d = [L : K]$ a small prime. The form of the total dimension $N = d^2(d-1)2^{k-1}$ constrains our choice of dimension, but for examples of cryptographically relevant sizes with $d = 3$ one can consider setting $k = 6$ or $k = 7$ to achieve dimension $N = 576$ or $N = 1152$ respectively. If one required additional flexibility of dimension one could also consider increasing d or replacing the power-of-two cyclotomic field with any cyclotomic field whose intersection with $\mathbb{Q}(\zeta_d)$ is precisely $\mathbb{Q}$. This method comes with the subtle drawback that the module rank d is also present in the dimension of the base field K , which precludes the case where one wants a large module rank and a small center. On the other hand, since such cases are excluded in our security proof we view this drawback as minor.

4.4 Extensions Where q Splits Completely

All suggested algebras in the previous section satisfy the conditions required for our chosen norm $\|\sigma_{\mathcal{A}}(x)\|_2$ to be well-defined and our search-decision reduction to hold. In particular, they have root of unity non-norm γ and K is cyclotomic. Later in this thesis, in order to enable efficient multiplication algorithms, it will turn out that it is convenient to have a modulus q that

splits completely into a product of prime ideals in both $\mathcal{O}_K$ and $\mathcal{O}_L$. An appeal to Chebotarev's Density Theorem suggests that a proportion of $1/d$ of the primes q that split completely in K also do so in L . In cases where d is small this suggests that finding such primes should not prove too arduous, but since cryptosystems require specific parameters rather than density arguments we provide constructions satisfying the requisite conditions on q .

Because K is a cyclotomic field in the above constructions and any $q = 1 \pmod{m}$ splits completely in $\mathbb{Q}(\zeta_m)$ it is straightforward to find q which splits completely in $\mathcal{O}_K$. We would also like q to be of roughly appropriate cryptographic size (say between 3000 and 15000 as a soft estimate, once again presuming parameters similar to those of NewHope or Kyber). Having q split completely in L is not as straightforward as in K because L is not a cyclotomic field, so we return to our examination of the proof of Theorem 4.1. Recall that in this proof the extension field L is a subfield of $K(\zeta_s)$ for some prime integer s satisfying $s = 1 \pmod{n}$ and, for $n = p^a$, p^{a+1} does not divide $s - 1$. That is, n is the highest power of p that divides $s - 1$. We have several methods to ensure that q splits completely in L , of which we start with the most naive.

4.4.1 Naive Method

For our general method we rely on the following fact: If $\mathfrak{q}_i$ is an ideal of $\mathcal{O}_K$ which splits completely in an extension M/K then it splits completely in any intermediate field $M/L/K$. As it is conceptually simpler to apply this idea to the integer q than to the $\mathcal{O}_K$ -ideals $\mathfrak{q}_i$ we use a simpler statement, that if $\langle q \rangle$ splits completely in some M containing L then it splits completely in L . This gives us an easy way to find some q that splits completely by examining a cyclotomic field that contains L : let $K = \mathbb{Q}(\zeta_n)$ and let $M = K(\zeta_s)$. Then since $s = 1 \pmod{n}$ it follows that $M = \mathbb{Q}(\zeta_{sn})$. Thus q splits completely in M if and only if $q = 1 \pmod{sn}$ and consequentially splits completely in our extension L if $q = 1 \pmod{sn}$. Since there are infinitely many primes equal to $1 \pmod{sn}$ this recipe always provides a prime q that splits completely in L . The upside of this method is that it is both very general and simple, since all candidate fields L we construct are contained in a larger cyclotomic field. Theoretically, this method can be extended to any abelian extension of $\mathbb{Q}$ using the partial converse of the

Kronecker-Weber Theorem. However, using the Kronecker-Weber Theorem constructively is not as straightforward as picking s as in the proof of Theorem 4.1, so this extension to general abelian L is slightly contrived.

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

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

4.4.2 Quadratic Case

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

see that $s \equiv 1 \pmod{4}$ and that $7681 = 3788^2 \pmod{12289}$, so that $q = 12289$ splits completely in E, K , and thus L , as required. Since this prime is explicitly the prime used in NewHope for all parameter sets we view this method as a substantial improvement on the previous technique.

4.4.3 Quartic Fields

Unfortunately, the quartic case where $d = 4$ is a little harder to understand. This is because quartic subfields of cyclotomic fields do not permit such a simple and explicit description as their quadratic counterparts do. However, we can still make some progress in explicitly understanding the factorization of q in these objects using the method of describing L as a compositum MK/K . Now, M will be a quartic subfield of the field $\mathbb{Q}(\zeta_s)$ and one can establish the linearly disjoint nature of M and K required to express L as this compositum by e.g. examining their discriminants: since K is a power-of-two cyclotomic field the only prime appearing in its discriminant is 2, and since M is a subfield of $\mathbb{Q}(\zeta_s)$ the only prime in its discriminant is s . Since they have coprime discriminants they are linearly disjoint, and since ramified primes are factors of the discriminant we have a relatively easy way to discount q being ramified ($q \neq 2, s$), so the remaining case to concern ourselves with is q being inert.

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

In fact, this quartic method can be applied to other instances where we do not have an explicit description of the subfields of $K(\zeta_s)$ which have degree d over K : define the families of q which split completely in K , then check whether those q split completely in L using computational

software. Since $q = 1 \pmod m$ and m is relatively large, there will not be many q to check of appropriate size for lattice cryptography, and so we conclude that this method is sufficient for fixed choices of fields L, K for which a satisfactory q exists.

Chapter 5

CLWE in Cryptography

The Force is with you, young
Skywalker, but you are not a Jedi yet.

Darth Vader, *Star Wars V: The Empire
Strikes Back* By George Lucas

In the previous two chapters we demonstrated that the hardness of the search CLWE problem derives from the assumed hardness of structured short vector problems and that the decision CLWE problem is as hard as the search variant with restricted secrets. We then presented appropriate families of algebras within which this derivation holds. We have also shown that concrete algebras within these families exist in dimensions suitable for cryptography. We are nearly ready to apply the CLWE problem to cryptography, but first we must demonstrate how to mold the CLWE assumption into a structure more familiar to lattice cryptosystems. Namely, we will show the hardness of CLWE in both normal form and primal form, presenting the CLWE problem as a discrete problem with small secrets. Following this we will present a cryptosystem in the form of the example scheme in Section 2.7 that is based on the CLWE problem. Finally, we present some asymptotically fast algorithms for computing multiplication in natural orders modulo q , the most costly operation required in the cryptosystem, and demonstrate the weaknesses of poorly chosen γ when CLWE is in primal form.

5.1 Discrete CLWE with Small Secrets

5.1.1 Primal Form

The primal form of an LWE problem is what allows for the conceptually simple discrete representation. In the primal form, a sample from the RLWE distribution is taken over the discrete space $R_q \times R_q$ rather than $R_q \times K_{\mathbb{R}}/R_q^{\vee}$ at no loss of security. We will show that the same transformation can be performed to define primal CLWE, whose samples will be in the analogous space $\Lambda_q \times \Lambda_q$. We begin with an adaptation of Lemma 2.23 of [LPR13] which discretizes the error distribution to $\Lambda_q^{\vee}$. In the following lemma we have followed the example of [LPR13] and remodelled CLWE samples by replacing $\Lambda^{\vee}$ with $q\Lambda^{\vee}$ and scaling up b by a factor of q . Since q is central this change is purely semantic, and is done to permit more legible arguments.

Lemma 5.1. *Let $\lfloor \cdot \rfloor$ be a valid discretization from $\bigoplus_{i=0}^{d-1} u^i \mathcal{O}_{L_q}^{\vee}$ to $q \cdot \Lambda_q^{\vee}$. Then, there exists an efficient transformation which on input a pair $(a', b') \in \Lambda_q \times (\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}})/q\Lambda_q^{\vee}$ outputs a pair $(a, b) \in \Lambda_q \times \Lambda_q^{\vee}$. Furthermore, if the input pair (a', b') is sampled from the uniform distribution then the output of this transformation is uniformly distributed and if $(a', b') \leftarrow \Pi_{q,s,\psi}$ is distributed according to the CLWE distribution for secret s and error distribution ψ then the output is distributed according to $\Pi_{q,s,\chi}$ for $\chi = \lfloor \psi \rfloor$, the distribution obtained by sampling from ψ and discretizing the output.*

We do not present a proof, since it is identical to that of Lemma 2.23 of [LPR13] up to a change of ambient space. Our statement is less general than theirs', which allows for both switching of the modulus q and changing the centre of the error distribution χ , but since we do not need these two functionalities we present the simpler statement above. Nonetheless, we observe that since the modulus is an integer, and thus central, we could generalize our statement in the same manner and the proof would still follow theirs'. Before we can demonstrate the hardness of the CLWE primal form we require a lemma replacing the dual $\Lambda_q^{\vee}$ with the natural order itself.

Lemma 5.2. *Let $\mathcal{A} = (L/K, \theta, \gamma)$ and let $\gamma \in \mathcal{O}_K$ be a unit. Then, there exists an efficient transformation which takes a sample $(a', b') \in \Lambda_q \times \Lambda_q^{\vee}$ and outputs a sample $(a, b) \in \Lambda_q \times \Lambda_q$.*

Furthermore, if the input pair (a', b') is sampled from the uniform distribution then the output of this transformation is uniformly distributed and if $(a', b') \leftarrow \Pi_{q, s', \chi'}$ as in Lemma 5.1 is distributed according to the CLWE distribution for secret s' and error distribution χ' then the output is distributed according to $\Pi_{q, s, \chi}$ for some error distribution χ over Λ_q and secret $s \in \Lambda_q$.

Proof. Recall that since γ is a unit, $\Lambda_q^\vee = \bigoplus_i u^i \mathcal{O}_{L_q}^\vee$. It is well known (see e.g. [Pei16]) that there is a tweak factor t which induces a bijection modulo q taking $\mathcal{O}_{L_q}^\vee$ to $\mathcal{O}_{L_q}$ by multiplying by t , the ring analog of our Lemma 3.7. The transformation takes the sample (a', b') and outputs $(a, b) = (a', b't)$. Since multiplying by t induces a bijection from $\mathcal{O}_{L_q}^\vee$ to $\mathcal{O}_{L_q}$, it is clear that our transformation maps uniform samples to uniform samples¹. Otherwise, $(a', b') \leftarrow \Pi_{q, s', \chi'}$, and so the resulting pair satisfies $(a, b) = (a', a's't + et)$. Setting $s = s't \in \Lambda_q$ and letting χ be the distribution over Λ_q derived by sampling from χ' then multiplying on the right by t completes the proof. $\square$

Combining these two lemmas provides us with the following argument: since there exists a valid discretization preserving the uniform distribution and mapping $\Pi_{q, s, \psi}$ to its discretization $\Pi_{q, s, \chi'}$, the problem of distinguishing $\Pi_{q, s, \chi'}$ from uniform is at least as hard as the decision CLWE problem as presented in Section 3.1. Then, since there is a further transform preserving the uniform distribution and mapping $\Pi_{q, s, \chi'}$ to the primal form distribution $\Pi_{q, s, \chi}$ over $\Lambda_q \times \Lambda_q$, the primal CLWE decision problem is also at least as hard as the dual CLWE decision problem. We state the primal distribution formally.

Definition 5.3 (Primal CLWE Distribution). Let $\mathcal{A} = (L/K, \theta, \gamma)$ be a cyclic division algebra with unit γ satisfying $|\gamma| = 1$ and let Λ denote its natural order. For prime $q \geq 2$, $s \in \Lambda_q$, and an error distribution χ over Λ_q , the primal form $\Pi_{q, s, \chi}$ of the CLWE distribution is sampled from by sampling $a \leftarrow \Lambda_q$ uniformly at random, sampling $e \leftarrow \chi$, and outputting $(a, b) = (a, a \cdot s + e) \in \Lambda_q \times \Lambda_q$. The primal form distribution permits the search and decision CLWE problems as usual.

Remark. The tweak factor of t from the proof of Lemma 5.2 will effect the error distribution one should use for a practical CLWE cryptosystem. In the case of RLWE in a 2-power cyclotomic

¹Note that t must be applied on the right since it is not necessarily central.

field the tweak factor is precisely n , the dimension of the underlying number field. Since our field L will not be cyclotomic, t may not take such a convenient form. However, we could equivalently replace this t with an element $t \in \mathcal{O}_K$ obtained as in Lemma 3.6 which will typically be the tweak factor of the center K . Since we propose field's whose centers are the fields used in RLWE, this tweaking has the same effect on our error as it does on error terms in RLWE. Furthermore, this t is non-unique, and we could instead have adapted the more involved techniques of [RSW18] to select a t in the ring with short entries.

5.1.2 Normal Form

We now show that the primal form of the CLWE distribution permits small secrets instead of uniform ones. Precisely, we will show that primal CLWE is still hard in its normal form, in which the secret s is sampled from the error distribution χ . We require two facts for our proof: firstly, given that q splits completely in K the ring Λ_q is isomorphic to the direct product of n full matrix algebras over $M_{d \times d}(\mathbb{F}_q)$, which can be seen by appealing to the CRT-style decomposition of Lemma 3.17 and Wedderburn's Theorem as in Section 3.1, Propositions 1 and 4 of [OAS12]. Secondly, we require that a non-negligible fraction in n of elements of Λ_q are invertible, which follows for fixed, small, d and $q \in \text{poly}(n)$ from this direct product decomposition. Otherwise, our proof follows the outline for that of plain LWE from [ACPS09]. Given these two facts, we proceed with showing that the normal form of the CLWE distribution is as hard as the case of taking the secret uniformly at random.

Lemma 5.4. *For a fixed d and $q \geq (n + 1)$, a non-negligible proportion of elements of Λ_q are invertible.*

Proof. Following the decomposition of Lemma 3.17 and Wedderburn's Theorem, it is sufficient to show that a non-negligible proportion of elements of

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

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

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

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

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

Now that we know that a sufficient proportion of elements of Λ_1 are invertible, we are ready to demonstrate security of CLWE in the normal form.

Lemma 5.5. *There is a probabilistic polynomial time reduction from the CLWE problem with uniformly random secret s and error distribution χ to the CLWE problem with secret $s' \leftarrow \chi$.*

Proof. It is sufficient to show that there is an efficient transformation that preserves the uniform distribution and maps samples with secret s to samples with some new secret s' taken from χ . To do so, we sample pairs $(a, b) \leftarrow \Pi_{q,s,\chi}$ until a pair $(a_1, b_1 := a_1 \cdot s + e_1)$ such that a_1 is invertible in Λ_q is obtained. Since a non-negligible fraction of elements of Λ_q are invertible by Lemma 5.4, this step runs in probabilistic polynomial time.

Now, given a pair (a_i, b_i) , we obtain a new sample by outputting $(\bar{a}_i, \bar{b}_i) = (a_i a_1^{-1}, a_i a_1^{-1} b_1 - b_i)$. Since a_1^{-1} is invertible, $\bar{a}_i$ is uniform and it is clear that on input a uniform sample we output

a uniform sample. Similarly on input (a_i, b_i) from $\Pi_{q,s,\chi}$

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

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

5.2 Example CLWE Cryptosystem

We are now ready to present an example cryptosystem based on the hardness of the CLWE problem. Our public key scheme will be the natural CLWE variant of that of Section 2.7, and as required for schemes in this form we will use the primal and normal variant of the CLWE problem. We require the following parameters:

- An algebra $\mathcal{A} := (L/K, \theta, \gamma)$ satisfying the conditions for the security reductions of Section 3.6.
- A prime modulus $q = 1 \pmod m$, where K is the m^{th} cyclotomic field.
- An error distribution χ over Λ_q .

We will denote with bold faced letters the vector form of an element of Λ_q , e.g. if $a = a_0 + ua_1 + \dots + u^{d-1}a_{d-1}$ then $\mathbf{a} = (a_0, a_1, \dots, a_{d-1})$. We note that $\mathcal{O}_{L_q}$ has a polynomial representation of dimension $n \cdot d$, and so we encode our message $m \in \{0, 1\}^{n \cdot d^2}$ as an entry of Λ_q as a vector $\mathbf{m}$ of length d whose entries are each $\{0, 1\}$ polynomials of degree $nd - 1$. The steps of the scheme are as follows:

KeyGen: Alice generates a CLWE sample $(a, b := a \cdot s + e)$, where $a \in \Lambda_q$ is uniformly random and $s, e \leftarrow \chi$, and outputs the public key $\mathbf{a}, \mathbf{b}$ in vector form.

Encrypt: To encrypt $m \in \{0, 1\}^{n \cdot d^2}$, Bob samples $t, e_1, e_2 \leftarrow \chi$ and outputs the ciphertext pair $\mathbf{u}, \mathbf{v}$ where $\mathbf{u} := \phi(a)^T \cdot \mathbf{t} + \mathbf{e}_1, \mathbf{v} := \phi(b)^T \cdot \mathbf{t} + \mathbf{e}_2 + \lfloor \frac{q}{2} \rfloor \cdot \mathbf{m}$.

Decrypt: To decrypt, Alice computes $\mathbf{c} = \mathbf{v} - \phi(s)^T \cdot \mathbf{u}$ and recovers m by rounding the entries of each constituent polynomial of $\mathbf{c} \in \mathcal{O}_{L_q}^d$ separately to 0 or $\lfloor \frac{q}{2} \rfloor$ and outputting 0 or 1 respectively.

We observe two benefits of instantiating this scheme in the cyclic algebra setting rather than over modules as in [BDK⁺18], 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 the embedding $\phi(a)$ lets us generate a matrix and thus saves a factor of d in the size of the public key². Secondly, by extending $\mathbf{b}$ to $\phi(b)$ we are able to increase the dimension of $\mathbf{v}$, and correspondingly increase the size of the message, by a factor of d . In this way our construction preserves dimensions, which may be of use either in the case where one wants to compose operations to build cryptosystems which have more complex functionalities than a simple public key encryption scheme, or in the case where one wanted to apply an error correcting code on a shorter message m .

We will shortly demonstrate security and correctness of the scheme, but first 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 $\phi(a)^T \cdot \mathbf{s} + \mathbf{e}$ is a valid CLWE instance in some other algebra $\mathcal{A}'$ for which the CLWE problem is still hard.

Lemma 5.6. *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 left regular representation $\phi'(a')$ and natural order Λ' such that for $a \in \Lambda$ there exists $a' \in \Lambda'$ satisfying $\phi(a)^T = \phi'(a')$. Moreover, $\mathcal{A}'$ is a division algebra 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 and thus closed under inverses. Λ'_q and Λ_q

²Of course, for a public key scheme Alice may circumvent this problem in practice by publishing a seed to generate $\mathbf{A}$ using a PRG, but there are more complicated applications where this advantage may be realised.

are additively isomorphic because both algebras share the same underlying fields and γ, γ^{-1} are units of $\mathcal{O}_L$. Since the first row of $\phi(a)$ is precisely $(a_0, \gamma\theta(a_{d-1}), \gamma\theta^2(a_{d-2}), \dots, \gamma\theta^{d-1}(a_1))$, by setting $a' = a_0 + u\gamma\theta(a_{d-1}) + \dots + u^{d-1}\gamma\theta^{d-1}(a_1)$ and observing that θ fixes γ and θ^d is the identity it is easy to check that $\phi(a)^T = \phi'(a')$. $\square$

The proofs of correctness and security of the cryptosystem are similar in spirit to those of other structured LWE schemes such as e.g. NewHope [ADPS16] or Kyber [BDK⁺18]. We proceed with a security argument.

Lemma 5.7. *The defined scheme is IND-CPA secure under the assumption that the decision CLWE problem given access to few samples in the normal form is hard*

Proof. The goal of an IND-CPA adversary is to distinguish, with non-negligible advantage, between encryptions of two plaintexts m_0, m_1 . The challenger chooses $i \in \{0, 1\}$ uniformly at random and encrypts m_i as $\mathbf{u}_i, \mathbf{v}_i$. By the assumption that the decision CLWE problem is hard, the adversary cannot distinguish between the case where the public element is correctly formed as $b = as + e$ and the case where it is replaced by a uniformly random b' , so we replace the challenge ciphertext $\mathbf{v}_i$ with $\mathbf{v}'_i$ by replacing b with b' . Setting $\mathbf{v}''_i := \mathbf{v}'_i - \lfloor \frac{q}{2} \rfloor \cdot \mathbf{m}_i$, it follows by Lemma 5.6 that $\mathbf{u}_i, \mathbf{v}'_i$ represent two samples from a valid CLWE distribution with small secret $\mathbf{t}$, and so the adversary cannot distinguish them from uniform with non-negligible advantage. Hence, the challenger cannot distinguish $\mathbf{v}''_i$ from uniform with non-negligible advantage. Following the chain of replacements it follows that the challenger can not distinguish $\mathbf{u}_i, \mathbf{v}_i$ from uniform and thus can not predict i with non-negligible advantage. $\square$

This proof is essentially the same as in the ring and module cases, arguing that an attacker can not derive any information from the observable information $a, b, \mathbf{u}, \mathbf{v}$ since they represent a collection of CLWE samples under appropriate secrets. We now move on to showing conditions under which the scheme is correct and Alice obtains the intended message m .

Lemma 5.8. *The defined scheme is correct as long as the ℓ_∞ norm of $\mathbf{e}' = (\phi(e)^T \cdot \mathbf{t} + \mathbf{e}_2 - \phi(s)^T \cdot \mathbf{e}_1)$ is less than $\lfloor \frac{q}{4} \rfloor$, where the ℓ_∞ norm is over the vector of all polynomial coefficients of each u^i entry of $\mathbf{e}'$ of dimension $n \cdot d^2$.*

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

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

from which the result follows immediately. $\square$

We note that the error term $\mathbf{e}'$ will be unsurprising to those familiar with other forms of LWE based cryptography. The error parameters for our decision reduction are equivalent to those of MLWE up to some sparse, small, covariance terms. Thus, one could consider the marginal distributions and appeal to e.g. the proof of correctness with high probability of Kyber to show the correctness of each of the d constituent polynomials of m independently, and then argue that since we are intending to keep d small the union bound proves correctness of our scheme with high probability. In practice, since we support a message m with dimension a factor of d larger than that of the corresponding MLWE schemes but the same dimension as that of the ring variant, we would recommend using some form of error correcting code (ECC) common to ring schemes such as the repetition code used in NewHope. The use of an ECC would require more robust error analysis, but would strictly improve the correctness of the scheme at a corresponding cost in message length.

5.3 Operational Complexity in Cyclic Algebras

In the previous section we showed that the CLWE problem can be used to construct a standard LWE based cryptosystem. Assuming that parameters across all variants of the LWE assumption are roughly equivalent, the CLWE problem supports key and message sizes as advantageous as those of the RLWE problem, and better than those of the module case. Along with storage considerations, another important facet of the ambient space in LWE cryptography

is the efficiency of operations. Here, we will construct algorithms and consider the asymptotic complexity of multiplication in a cyclic algebra in order to compare it to the ring and module variants. Since in practice we consider operations modulo some prime q , addition in rings, modules, and cyclic algebras can be considered as addition in vector spaces over $\mathbb{Z}_q$, which has complexity dominated by that of multiplication.

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

- In the ring case, the operation $\mathbf{A}\mathbf{s}$ over $\mathbb{Z}_q$ is a representation of the ring operation $a \cdot s$ in $R_q \cong \mathbb{Z}_q[X]/(X^N + 1)$. Using the CRT decomposition in dimension N of [LPR13], this operation is decomposed into coordinatewise multiplication in a vector of dimension N over $\mathbb{Z}_q$, following which the decomposition is reversed to recover $a \cdot s$. The complexity of this technique is dominated by that of the CRT decomposition, which takes time $O(N \log N)$, although the coordinatewise multiplication also requires time $O(N)$.
- In the module case, $\mathbf{A}$ is a $d \times d$ matrix over R_q . In this case, one can compute $\mathbf{A}\mathbf{s}$ by applying the CRT in dimension n coordinatewise on $\mathbf{A}$ and $\mathbf{s}$. This requires $d^2 + d$ applications of the CRT, for a total asymptotic complexity of $O(d^2 n \log n) = O(Nd \log(N/d))$. Again, this hides a coordinatewise multiplication step which takes time $O(Nd)$ in this setting.
- In the cyclic algebra case, $\mathbf{A}$ is a matrix in the shape $\phi(a)$, where $\phi(a)$ is the left regular representation of $a \in \Lambda_q$. We will compute the complexity of the operation $\phi(a) \cdot \mathbf{s}$ in this section. Explicitly, our algorithm has complexity $O(N \log N/d^2) + O(Nd^{\omega-2})$ in the case where q splits completely in L , with ω denoting the exponent of matrix multiplication.

The latter term corresponds to the cost of multiplication in our analog of the finite fields used in the CRT method for RLWE. However, we are forced to store elements in an atypical manner to permit our techniques.

We note that the complexities of the ring and cyclic algebra cases are broadly equivalent since the small difference in the logarithm term is relatively marginal compared to the matching $O(N)$ parts, whereas both outperform the module case by a factor of d . Asymptotically, the complexity of our CLWE multiplication algorithms do outperform those of techniques for the ring case, but since we will be forced to store elements in a different manner than is typical for cryptosystems using the ring case we recommend caution when comparing these operations.

5.3.1 Algorithm for Multiplication in Cyclic Algebras

We recall some details necessary to understand our multiplication algorithm. Recall that in the explicit constructions of Theorem 4.3 the base field K is cyclotomic and q is a prime integer chosen so that $\langle q \rangle$ splits completely in $\mathcal{O}_K$ as $\langle q \rangle = \mathfrak{q}_1 \dots \mathfrak{q}_n$, where n is the dimension of K as an extension of $\mathbb{Q}$. Furthermore, the degree of L over K is a typically small d . Then, following the CRT-like decomposition of Lemma 2.34 we write

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

for $\mathcal{R}_i = \bigoplus_{j=0}^{d-1} u^j \mathcal{O}_L / \mathfrak{q}_i \mathcal{O}_L$. We will show that each $\mathcal{R}_i$ is a skew polynomial ring over $\mathbb{Z}_q$, and in particular a skew polynomial ring for which we can apply the algorithms of [CLB17] to compute multiplication independently in each $\mathcal{R}_i$ in $O(d^\omega)$ operations in $\mathbb{Z}_q$, which output elements whose u coordinates are in the form $\sum_i \ell_i k_i$ for $k_i \in \mathcal{O}_{K_q}$ and $(\ell_i)_i$ some arbitrary normal basis for $\mathcal{O}_{L_q}$ over $\mathcal{O}_{K_q}$. We remark that the representation as a skew polynomial ring need not contradict the fact that we viewed the rings $\mathcal{R}_i$ as matrix rings in Theorem 3.20, since computing matrix multiplication can be reduced to the problem of computing multiplication of skew polynomials (see [CLB17]). Since $\omega \leq 3$ and we can compute the multiplication in each $\mathcal{R}_i$ in parallel, this leads to a complexity of approximately $O(Nd)$. However, we must also

compute the complexity of the splitting isomorphism, which we will do in the next section.




Figure 5.1: Depiction of Multiplication Algorithm.

5.3.2 The Rings $\mathcal{R}_i$

In order to apply the algorithm of [CLB17], we must confirm that each $\mathcal{R}_i$ satisfies the following conditions:

- $\mathcal{R}_i$ is the quotient of a skew polynomial ring with center $\mathcal{O}_K/\mathfrak{q}_i$ by a polynomial in the form $X^d - \gamma$.
- γ is a norm from $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ into $\mathcal{O}_K/\mathfrak{q}_i$.³
- $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ is a field extension of $\mathcal{O}_K/\mathfrak{q}_i$ or an étale- $\mathcal{O}_K/\mathfrak{q}_i$ algebra.⁴

The first of the conditions follows immediately from the definitions of a skew polynomial ring and a cyclic algebra. The veracity of the latter conditions will depend on how the prime ideal

³Due to the modulo reduction this does not contradict the assumption that γ is not a global norm.

⁴For the purposes of this thesis, it is sufficient to understand an étale- $\mathcal{O}_K/\mathfrak{q}_i$ algebra to be a direct product of finitely many copies of $\mathcal{O}_K/\mathfrak{q}_i$; the formal definition is somewhat convoluted.

$\mathfrak{q}_i$ of $\mathcal{O}_K$ splits in $\mathcal{O}_L$ as $\mathfrak{q}_i\mathcal{O}_L$. Since $\mathfrak{q}_i$ is prime in K and L/K is Galois, we know

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

for some prime ideals $\mathfrak{q}_{i,j}$ in $\mathcal{O}_L$ and integers e, g satisfying $efg = [L : K] = d$, where f denotes the inertial degree. Assuming that L is constructed as a subfield of a cyclotomic field as in [LMM08], it is a Galois number field and it follows that each $\mathfrak{q}_i$ splits with the same e, f , and g . Furthermore, since they are coprime as ideals of $\mathcal{O}_K$, their factorizations' in L are disjoint. Thus, we are left to consider three cases.

We first consider the case where each $\mathfrak{q}_i\mathcal{O}_L$ remains prime in $\mathcal{O}_L$. It follows that $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ is a finite field, and computing the norm of $\mathfrak{q}_i\mathcal{O}_L$ indicates $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L \cong \mathbb{F}_{q^d}$. In this case it is easy to see that $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ is a finite field extension of $\mathcal{O}_K/\mathfrak{q}_i \cong \mathbb{F}_q$ and consequentially, because the norm map is surjective over finite field extensions, that γ is a norm. Here it is clear that the algorithms of [CLB17] can be applied.

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

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

where each $\mathcal{O}_L/\mathfrak{q}_{i,j} \cong \mathbb{F}_q$, and it follows that $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ is an étale- $\mathcal{O}_K/\mathfrak{q}_i$ algebra. We are left to show that γ is a norm, which we show via the stronger condition that the norm map in this extension is surjective. By the CRT, $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ is isomorphic to a direct product of d copies of $\mathbb{F}_q$. Since the embeddings of L cyclically permute the ideal factors of $\mathfrak{q}_i$ it follows that the relative norm of an element $(x_1, \dots, x_d) \in \bigotimes_{j=1}^d \mathcal{O}_L/\mathfrak{q}_{i,j}$ is precisely $\prod_{k=1}^d x_k \pmod{q}$. It is easy to see that this norm is surjective (because any $x \in \mathbb{F}_q$ is the norm of e.g. $(1, 1, \dots, x)$) and now once again we can apply the multiplication algorithms of [CLB17].

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

The final case to consider is the ramified case, when $e \neq 1$. Now the factorization of $\mathfrak{q}_i \mathcal{O}_L$ contains some power $\mathfrak{p}_i^{e_i}$ of a prime $\mathcal{O}_L$ ideal $\mathfrak{p}_i$. In this case, we are not able to verify that the necessary conditions for the algorithms of [CLB17] hold. However, we observe that the ideal $\langle q \rangle$ ramifies in $\mathcal{O}_L$ if and only if q divides the discriminant of $\mathcal{O}_L$. Since only a finite number of primes divide this discriminant, we restrict ourselves to considering the cases where q does not ramify. We emphasize that in the main cases of interest, where K is the m^{th} cyclotomic field with m having small divisors and $[L : K]$ is small, it is particularly unlikely that the large modulus q typical in cryptography divides the discriminant of L . Indeed, when we pick L as a subfield of $K(\zeta_s)$ for some large prime integer s using the techniques of [LMM08] as in Theorem 4.3, it is easy to quantify which primes potentially ramify for a fixed choice of fields: either s or the primes smaller than or equal to the divisors of m . As an easy example, the modulus $q = 12289$ does not ramify in the example algebras given in the Section 4.3 achieving dimension 1024.

5.3.3 Complexity of the CRT Style Isomorphism

We have shown that we may apply the algorithms of [CLB17] to compute the multiplication operation in each $\mathcal{R}_i$ in complexity $O(d^\omega)$. We are left to consider the complexity of the isomorphism defined by Lemma 2.34 generating the rings $\mathcal{R}_i$. Essentially, this operation is a coordinatewise split of the u coordinates of $\Lambda_q = \bigoplus_{j=0}^{d-1} u^j \mathcal{O}_L$, where each entry is split into its mod $\mathfrak{q}_i \mathcal{O}_L$ parts. That is, the isomorphism maps

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

Splitting one element $x_i \in \mathcal{O}_K$ can be done in time $O(n \log n)$ using the CRT algorithm of [LPR13] when K is a cyclotomic field of dimension n . However, L is not a cyclotomic field, but instead a small degree d cyclic extension of a cyclotomic. Furthermore, we are trying to split the elements of L modulo ideals of K extended to those of L . We do not know of an existing general, efficient way of doing this. The naive estimate for an optimal method would

take time $O(nd \log nd)$, where nd is the dimension of L , but we suspect something this efficient is impossible. We have to perform d such splits, which would result in a total complexity of $O(N \log N/d)$. Note that this compares relatively closely with the $O(Nd^{0.3})$ claimed for the multiplication step, and since these steps are sequential rather than parallel which of them dominates the asymptotic complexity would depend on the exact relationship between n and d , but the result is an operational complexity essentially equivalent to that of the ring variant.

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

5.3.4 Fast Cryptography When q Splits Completely in L

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

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

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

and similarly

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

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

We exhibit this with a trivial example: set $d = 3, n = 2$. Then the action of θ on $(\mathbb{F}_q^3)^2$ is

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

A valid $\mathcal{O}_K/\mathfrak{q}_i$ basis for $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ of size d is $\mathbf{e}_1, \dots, \mathbf{e}_d$, where $\mathbf{e}_i = (0, \dots, 1, \dots, 0)$ denotes the i^{th} element of the standard basis of dimension d . Furthermore, this basis is orthonormal in the sense that $\mathbf{e}_i \cdot \mathbf{e}_j = \mathbf{e}_i$ for $i = j$ and 0 otherwise and cyclic⁵ in the sense that $\theta(\mathbf{e}_i) = \mathbf{e}_{i+1}$ (e.g. normal), since the Galois group $\langle \theta \rangle$ of L over K permutes the factors $\mathfrak{q}_{i,j}$ of $\mathfrak{q}_i\mathcal{O}_L$ for each i . Because the CRT splits $\mathcal{O}_{Lq}$ into a direct product within which operations are computed coordinatewise, we can extend this to a basis of $\mathcal{O}_{Lq}$ over $\mathcal{O}_{Kq}$ in the finite field representation by concatenating n copies of this basis together, denoting by $\mathbf{e}_i^n$ the vector of dimension nd $(\mathbf{e}_i, \mathbf{e}_i, \dots, \mathbf{e}_i)$. This basis is still cyclic, with θ operating independently on each of the n copies of $\mathbb{F}_q^d$ and hence the n copies of $\mathbf{e}_i$. Concatenating the bases in this way also preserves the

⁵As long as you choose the ordering in the right way.

orthonormal property.

Denote the above basis by $\ell_1, \dots, \ell_d$. Recall that the CRT-like decomposition Lemma 2.34 splits each u coordinate, an element of $\mathcal{O}_{L_q}$, into its mod $\mathfrak{q}_i \mathcal{O}_L$ parts. However, we already know the mod $\mathfrak{q}_i \mathcal{O}_L$ parts of each ℓ_j by construction. So, if we store elements of $\mathcal{O}_{L_q}$ as $\ell = \sum_{j=1}^d \ell_j k_j$ for $k_j \in \mathcal{O}_{K_q}$ we can split ℓ into its $\mathcal{O}_L / \mathfrak{q}_i \mathcal{O}_L$ components in time $O(d \cdot n \log n)$ as long as the k_j elements are stored in the polynomial representation of $\mathcal{O}_{K_q}$. Consequentially, we can perform the CRT style decomposition of an element in Λ_q whose u coordinates are stored in this manner in time $O(d^2 \cdot n \log n) = O(N \log N / d^2)$.

Now we see a way to achieve fast multiplication in Λ_q . We are required to perform the CRT in each of the d many u coordinates, after which we can plug the rings $\mathcal{R}_i$ into the fast multiplication algorithm of [CLB17]. Since the CRT is an isomorphism and we know the image of ℓ_i under the CRT, this reduces to d copies of the CRT in $\mathcal{O}_K$, each with complexity $O(dn \log n)$, and therefore a total multiplication complexity of $O(N \log N / d^2) + O(Nd^{\omega-2})$ (where the latter term is $O(Nd)$ with naive matrix multiplication), which compares favourably with both ring and module multiplication in the same dimension. However, this algorithm comes with complications associated with the chosen representation of elements of $\mathcal{O}_{L_q}$, which we handle below.

Handling Elements in the Representation

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

We now discuss how to efficiently sample elements of Λ_q according to an appropriate error distribution. Recall from the security reduction of Lemma 3.16 that the error distributions we

recommend in practice are spherical or elliptical Gaussians in the coordinates of the embedding $\sigma_{\mathcal{A}}$. We sample using the following result.

Theorem 5.9. *Let L/K be a tower of number fields with $[K : \mathbb{Q}] = n$ and $[L : K] = d$ where K is a prime-power cyclotomic field. Let $q \geq 2$ be a prime modulus which splits completely in $\mathcal{O}_L$ and let $\ell_1, \dots, \ell_d$ be the cyclic basis of $\mathcal{O}_{L_q}$ over $\mathcal{O}_{K_q}$ satisfying $\ell_i \cdot \ell_j = \ell_i$ if $i = j$ and 0 otherwise. Then, the distribution on $\mathcal{O}_{L_q}$ obtained by sampling $k_1, \dots, k_d$ independently from a discrete Gaussian over $\mathcal{O}_{K_q}$ in the polynomial representation and outputting $\ell = \sum_i \ell_i k_i$ is a discrete Gaussian over $\mathcal{O}_{L_q}$ in the ℓ_2 norm over $L_{\mathbb{R}}$.*

Proof. Recall that in the case where K is a prime power cyclotomic the power basis is a rotation and a scaling of the canonical basis (see e.g. [CP16]), so a discrete Gaussian in the polynomial representation corresponds to a discrete Gaussian in the canonical basis as well. Order the canonical embedding of $\mathcal{O}_L$ such that elements of $\mathcal{O}_K$ embed as vectors of n blocks of length d that are the same in each block, e.g.

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

where each entry $k_{i,j}$ of k_i appears d times. Since the ℓ_i form a cyclic basis, in each d -block the entries of ℓ_{i+1} are just a cyclic shift of those of ℓ_i ⁶. For a fixed choice of basis the distribution in each d -block of ℓ is independent, because the $k_{i,j}$ are sampled independently from a spherical Gaussian. So we can consider one d block of ℓ at a time, and write the d -block of ℓ_1 as $a_1, \dots, a_d$. Since multiplication in the canonical embedding is coordinatewise and the ℓ_i form a cyclic basis, the first block of ℓ can be written as

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

Call the left matrix $\mathbf{A}$ and the right vector $\mathbf{k}$. The vector $\mathbf{k}$ is a Gaussian of parameter r , so $\mathbf{A}\mathbf{k}$

⁶Again assuming a sensible ordering.

has a Gaussian distribution with covariance matrix $r \cdot \mathbf{A}\mathbf{A}^T$ by e.g. Lemma 2.5 of [LVL18], and if this is diagonal and constant on the lead diagonal then we are done. Due to the structure of the canonical embedding and how we picked our basis in the $\mathcal{O}_L/\langle q \rangle$ representation, we have that $a_i = \theta^i(a_1)$ and that for $i \neq j$ $\theta^i(a_1) \cdot \theta^j(a_1) = 0 \pmod q$. It follows that the off-diagonal entries of $\mathbf{A}\mathbf{A}^T$ are 0 (since product being 0 is preserved under representations) and the diagonal entries are $\sum_{i=1}^d a_i$. Hence, the first d -block of ℓ is a spherical Gaussian distribution, and since this analysis holds for any block it follows that each block of ℓ is a spherical Gaussian. One also needs to show that the Gaussian distribution has the same variance in each block, but this follows from the fact that the K -embeddings permute the $\text{mod } \mathfrak{q}_i$ values and fix the ℓ_2 norm of $K_{\mathbb{R}}$. Explicitly, by construction each K embedding modulo $\langle q \rangle$ can be extended ‘identically’ onto $\mathcal{O}_L \pmod{\langle q \rangle}$ in a way that fixes each ℓ_i , so they must have the same set of values in each block (this would not be the case if we considered their norm in a global sense, and the restriction modulo q is strictly necessary). $\square$

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

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

Unfortunately, we are unable to generalize this theorem to the case where $\mathfrak{q}_i$ remains prime, or even intermediate cases. In this case, there exist cyclic bases of $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ over $\mathcal{O}_K/\mathfrak{q}_i$, but since $\mathcal{O}_L/\mathfrak{q}_i\mathcal{O}_L$ is a finite field and thus has no zero-divisors the cyclic bases are not orthogonal. Consequentially, the matrix $\mathbf{A}$ does not in general give a diagonal $\mathbf{A}\mathbf{A}^T$ and thus the distribution of $\mathbf{A}\mathbf{k}$ has several potentially large covariance terms. If one were able to tolerate the covariance,

the method can be extended in this case. It is also possible that a cyclic basis satisfying the condition that $\mathbf{A}\mathbf{A}^T$ is diagonal may exist for certain choices of field, but we were not able to find such a family of fields. We note that this question can be asked as a more generic question about finite fields: let $F = \mathbb{F}_{q^d}$ be a finite field with $d > 1$ and let θ denote the Frobenius automorphism of F . Does there exist a cyclic basis $b_1, \dots, b_d$ with $b_j = \theta^j(b_1)$ for F over $\mathbb{F}_q$ satisfying

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

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

Example 5.10. We exhibit an example of the basis ℓ_1, ℓ_2 in the simplest setting, that of a degree 2 extension of $\mathbb{Q}$. Let $L = \mathbb{Q}(i)$, with ring of integers $\mathcal{O}_L = \mathbb{Z}[i]$, and consider the ideal $\langle 5 \rangle$ of $\mathcal{O}_L$. Then 5 factorizes in $\mathcal{O}_L$ as $5 = (2+i)(2-i)$, and it is clear that $\langle 5 \rangle = \langle 2+i \rangle \cdot \langle 2-i \rangle$ is a decomposition into a product of prime ideals.

Using the notation $\mathbf{q}_1 := \langle 2+i \rangle, \mathbf{q}_2 := \langle 2-i \rangle$, it is easy to check that $2+i \equiv -1 \pmod{\mathbf{q}_2}$ and thus $-(2+i) = -2-i$ is a valid choice for ℓ_1 . Similarly, $-(2-i) = -2+i$ is an appropriate choice for ℓ_2 . Correspondingly, the distribution obtained by sampling $k_1, k_2 \leftarrow D_r$, the discrete Gaussian of parameter r over $\mathbb{Z}_5$, and outputting $k_1 \cdot (-2+i) + k_2 \cdot (-2-i)$ is a discrete Gaussian over $\mathcal{O}_L \pmod{\langle 5 \rangle}$ of parameter $|(-2+i) + (-2-i)| \cdot r = 4r$. Furthermore, to multiply two elements $k = k_1\ell_1 + k_2\ell_2$ and $g = g_1\ell_1 + g_2\ell_2$ modulo 5 one outputs $kg = (k_1g_1 \pmod{5}) \cdot \ell_1 + (k_2g_2 \pmod{5}) \cdot \ell_2$, at a cost of two operations in $\mathbb{Z}_5$, and performing the $\mathcal{O}_L \pmod{5}$ CRT on each u coordinate of an element of the resulting natural order Λ_5 can be done by merely reading off the $d^2 = 4$ values of k_i and no additional computation.

Furthermore, this is an example where the techniques of our next section may be advantageous. We will generalize the multiplication and CRT technique so that one is free to use any basis of $\mathcal{O}_L$ over $\mathbb{Z}$, for example the basis $\{1, i\}$. In this basis it is particularly easy to sample a discrete Gaussian in the polynomial representation of $\mathcal{O}_L \pmod{\langle 5 \rangle} \cong \mathbb{Z}_5[x]/(x^2 + 1)$, but the

resulting multiplication operation and CRT decomposition is not coordinatewise in the basis and so a small amount of efficiency is lost at a gain in parameter of the Gaussian. Specifically, to compute the CRT on an element $k = k_1 + k_2 \cdot i$, one has to precompute⁷ the values $i = -2 \bmod \mathfrak{q}_1, i = 2 \bmod \mathfrak{q}_2$ and output

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

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

5.3.5 Generalizing to non-Split q and Arbitrary Bases

In order to construct the cyclic, orthonormal, basis of Theorem 5.9, the previous section requires that q be completely split in both K and L . However, it is possible to drop the splitting condition in L and obtain fast multiplication algorithms in the general case at only a small loss of efficiency. We demonstrate the technique in this section and then briefly describe cases where a general algorithm may be superior to the one requiring that q splits by discussing alternatives to Theorem 5.9.

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

Choose an arbitrary integral $\mathcal{O}_K$ -basis $\ell_1, \dots, \ell_d$ of $\mathcal{O}_L$. As a precomputation phase, compute and store the images $\ell_j \bmod \mathfrak{q}_i \mathcal{O}_L$ for each i and j . The CRT-like decomposition of Lemma 2.34 splits each of the u coordinates of an element of Λ_q , an element of $\mathcal{O}_{Lq}$, into its $\bmod \mathfrak{q}_i \mathcal{O}_L$ parts. Once again, we suggest an algorithm where elements of $\mathcal{O}_{Lq}$ are stored in the form $\ell = \sum_{j=1}^d \ell_j k_j$ for $k_j \in \mathcal{O}_{Kq}$, e.g. on elements stored as K -combinations of this basis. We split $\ell \in \mathcal{O}_{Lq}$ into

⁷Note that precomputing the image of 1 is trivial.

its $\mathcal{O}_L/\mathfrak{q}_i$ components in time $O(d \cdot n \log n)$, since

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

where each $k_j \mod \mathfrak{q}_i$ can be computed in time $O(n \log n)$ by the K -CRT and each $\ell_j \mod \mathfrak{q}_i \mod \mathcal{O}_L$ was computed in the precomputation phase. Consequentially, we can perform the CRT style decomposition of an element in Λ_q whose u coordinates are all stored in this manner in time $O(d^2 \cdot n \log n)$, since we must split d^2 elements of $\mathcal{O}_K$. This decomposing complexity is the same as in the previous case where q splits completely. Following this, each ring $\mathcal{R}_i$ can be plugged in to the algorithm of [CLB17] to compute the multiplication in time $O(Nd^{\omega-2})$. However, since the ℓ_i do not correspond to a standard orthonormal basis we incur an extra cost when reversing this transformation. Namely, each of the u coordinates of each ring $\mathcal{R}_i$ is output by the algorithm of [CLB17] as an element $\ell \in \mathcal{O}_L \mod \mathfrak{q}_i \mathcal{O}_L$ expressed in an arbitrary normal basis. Before reversing the decomposition we must allow for the complexity of expressing each element of the output in the bases obtained by the images of $\ell_1, \dots, \ell_d \mod \mathfrak{q}_i \mathcal{O}_L$, as this basis was not necessarily normal. Since $\mathcal{O}_L \mod \mathfrak{q}_i \mathcal{O}_L$ is a vector space of dimension d over $\mathbb{F}_q$ this can be done via a precomputed change of basis matrix over $\mathbb{F}_q$ in time $O(d^\omega)$, and since there are n rings with d coordinates each the complexity of computing this on every coordinate is $O(nd^{\omega+1})$. The resulting multiplication algorithm has total complexity $O(N \log N/d^2) + O(Nd^{\omega-1})$. While this represents only a minor asymptotic loss, especially since we expect the first term to dominate the complexity, it is likely in practice that the extra step required to recover the basis representation would cause a tangible slowdown.

An unfortunate issue with this technique is that by replacing the orthonormal basis with an arbitrary basis we have lost Theorem 5.9 and thus the efficient method for sampling a discrete Gaussian in the representation $\ell = \sum_j \ell_j k_j$. However, this generalization allows for the use of an arbitrary basis $\ell_1, \dots, \ell_d$, unlike in the split case in which we chose a specific basis. Since we require that elements of Λ_q are input into the algorithm with u coordinates in the form $\sum_j \ell_j k_j$ this algorithm can be combined with the cryptosystem of Section 5.2 in the case where there is a basis $g_1, \dots, g_d$ of $\mathcal{O}_{Lq}$ over $\mathcal{O}_{Kq}$ in which one can compute the representation $\ell =$

$\sum_j g_j k_j$ particularly efficiently. This is because one can just sample ℓ from the usual Gaussian distribution over the polynomial basis of $\mathcal{O}_{Lq}$, compute its representation as $\ell = \sum_j g_j k_j$, and then apply the multiplication algorithm in this form. More generally, the flexible choice of basis allows for both non-split q and for a user to choose their favourite $\mathcal{O}_L$ basis properties, such as a normal basis or a basis consisting of small elements. We remark that it is likely possible to construct a pair of fields L/K that allow for a basis $\ell_1, \dots, \ell_d$ permitting a fast algorithm transforming from the polynomial representation of $\mathcal{O}_L$ to the representation $\sum_i \ell_i k_i$ with each k_i in polynomial representation, which would allow one to bypass the complications of sampling Gaussian distributions by just sampling in $\mathcal{O}_L$ directly.

5.3.6 Generalizing to Other Centers

In the exposition of the previous section we required that q splits completely in the center K . This corresponds to the requirement in the ring and module cases that q splits completely in the field K , which allows the use of the NTT to compute multiplications over a direct product of finite fields. However, there has been recent progress in loosening this requirement for the NTT and allowing the modulus q to be $1 \bmod n$ rather than $1 \bmod m$, where as usual K is the m^{th} cyclotomic field of degree n . For example, in the second round specification of Kyber [ABD⁺19] q is set as 3329 and $n = 256$, yet they still support efficient NTT based multiplication. In such cases, q is ‘well’ split but not completely split, and the fast NTT operations use the method of [LS19], where q splits into some product of prime ideals $\mathfrak{q}_i$ whose norms can be small powers of q .

We observe that our methods can be partially generalized to this case in the following manner. Say $\langle q \rangle = \prod_i \mathfrak{q}_i$ is a decomposition into prime ideals in $\mathcal{O}_K$ and there exists an efficient algorithm for fast multiplication in $\mathcal{O}_{Kq}$. We can replace our condition that q splits completely in $\mathcal{O}_L$ with the condition that each ideal $\mathfrak{q}_i$ in the $\mathcal{O}_K$ -factorization of q splits completely into a product of d prime ideals $\mathfrak{q}_i \mathcal{O}_L = \prod_{j=1}^d \mathfrak{q}_{i,j}$ in $\mathcal{O}_L$ of the same norm. Then, we can replicate the method of Section 5.3.4 to find a cyclic, orthonormal basis $\mathbf{e}_1, \dots, \mathbf{e}_d$ of $\mathcal{O}_L/\mathfrak{q}_i \mathcal{O}_L$ over $\mathcal{O}_K/\mathfrak{q}_i$ and concatenate together the bases for each i to make the cyclic, orthonormal, basis $\ell_1, \dots, \ell_d$

of $\mathcal{O}_{L_q}$ over $\mathcal{O}_{K_q}$. Since the basis is orthonormal, if $\ell = \sum_i \ell_i k_i$ and $g = \sum_i \ell_i g_i$ with each $k_i, g_i \in \mathcal{O}_{K_q}$, then

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

Since the basis is cyclic,

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

where we define $k_0 := k_d$.

Now we are able to use existing fast multiplication algorithms in $\mathcal{O}_{K_q}$ to compute operations in $\mathcal{O}_{L_q}$ by expressing elements in this basis. Represent each $x = \sum_{i=0}^{d-1} u^i x_i \in \Lambda_q$ by expressing each $x_i \in \mathcal{O}_{L_q}$ in the ℓ_j basis. Then, to multiply x and y in Λ_q one only has to compute multiplications in $\mathcal{O}_{K_q}$, since the operations required are just computing the non-commutative relation $\ell u = u\theta(\ell)$, which merely permutes the ℓ_i using θ , and computing multiplication and addition, which can be done coordinatewise in the orthonormal ℓ_i basis. Each L multiplication requires d multiplications in K , and each u coordinate of Λ requires d multiplications in L . Consequentially, naive multiplication in Λ_q takes d^3 instances of the efficient $\mathcal{O}_{K_q}$ -multiplication algorithm we have access to. For specific K -multiplication algorithms it is likely that this process can be streamlined; the intention of this section is merely to demonstrate that one can build efficient Λ_q operations from more general efficient operations over the center in the same manner that the techniques of Section 5.3.4 used the CRT method.

5.4 Attacking non-Division Algebras

In Section 3.3, the condition that γ is a non-norm element of L/K is required in order to stop parallelizing attacks in the style of that of [BCV18] applying to the CLWE problem. Thus,

γ is chosen so that γ^i is not in the norm group of L into K for $i = 1, 2, \dots, d-1$. Here, we demonstrate that picking γ that violates this condition leads to potentially vulnerable instances of the CLWE problem. We will need the following lemma, a rephrasing of Theorem 30.4 of [Rei75].

Lemma 5.11. *Let $\mathcal{A} = (L/K, \theta, \gamma)$ be a cyclic algebra with $[L : K] = d$. Let $\gamma, \delta \in K$ be non-zero. Then*

- $\mathcal{A} \cong \mathcal{A}_s := (L/K, \theta^s, \gamma^s)$ for each s such that $(s, d) = 1$.
- If $\gamma = 1$ then $\mathcal{A} \cong M_{d \times d}(K)$.
- If $\delta = N_{L/K}(\beta)\gamma$ for some non-zero $\beta \in L$ then $\mathcal{A} \cong \mathcal{A}' := (L/K, \theta, \delta)$.

Remark. If $\gamma \in \mathcal{O}_K$ is a unit then all isomorphisms of this lemma hold when replacing L and K with $\mathcal{O}_L$ and $\mathcal{O}_K$ respectively. The first and third can be seen by examining the proofs in [Rei75]; the first is a re-indexing of u coordinates of $\mathcal{A}$, and the third simply sends u to βu . The second requires a little more work. We map $\mathcal{A}$ to $\text{Hom}_K(L, L)$ by sending u to θ and $x \in L$ to the K -homomorphism on L defined by multiplying by x . Finally, we appeal to the standard isomorphism between $\text{Hom}_K(L, L)$ and $M_{d \times d}(K)$, which preserves integral elements as long as there exists an integral basis of L over K . We discuss the details of this last part later, because we also require it to preserve a notion of smallness.

Armed with this lemma, we demonstrate potential weaknesses of choosing γ poorly. Let $\mathcal{A} = (L/K, \theta, \gamma)$ be a cyclic algebra where γ lies in the norm group $N_{L/K}(L^\times)$ (and still lies in $\mathcal{O}_K$); later we will generalize our argument to the case where instead some power of γ less than d is a norm instead. Consider the primal CLWE instance $(a, a \cdot s + e) \in \Lambda_q \times \Lambda_q$, where a, s are uniform⁸ and $e \leftarrow \chi$ is drawn from an error distribution which is of Gaussian shape. Applying Lemma 5.11 transforms our sample into one over $M_{d \times d}(\mathcal{O}_{K_q}) \times M_{d \times d}(\mathcal{O}_{K_q})$. That is, we construct a sample in the form

$$(\mathbf{A}, \mathbf{A} \cdot \mathbf{S} + \mathbf{E})$$

⁸In practice s is typically sampled from the error distribution, but this normal form variant is no easier than the case where s is uniform. We assume uniform s here for ease of illustration

where $\mathbf{A}, \mathbf{S}, \mathbf{E} \in M_{d \times d}(\mathcal{O}_{K_q})$. Since isomorphisms are bijections, $\mathbf{A}$ and $\mathbf{S}$ are uniformly random matrices. Assume for the time being that the isomorphisms are also smallness preserving, so that if e is a small element of Λ_q then the corresponding matrix $\mathbf{E}$ will have entries that are small elements of $\mathcal{O}_K$.

Let $\mathbf{s}_i, \mathbf{e}_i$ denote the i^{th} columns of $\mathbf{S}$ and $\mathbf{E}$ respectively. Then, for each i the pair $(\mathbf{A}, \mathbf{A}\mathbf{s}_i + \mathbf{e}_i)$ constitutes d samples from the MLWE distribution in dimension n and rank d . That is, the single CLWE sample provides a collection of d samples from d instances of the MLWE distribution with different secrets $\mathbf{s}_1, \dots, \mathbf{s}_d$, where each set of samples shares the same uniformly random matrix $\mathbf{A}$. Since the difficulty of LWE problems is assumed to be superlinear in dimension N , solving d instances of the MLWE problem in dimension n and module rank d is easier than solving a single instance in dimension nd and rank d , the targeted dimension of our CLWE problem, which is essentially the parallelizing argument of the attack of [BCV18] on m -RLWE. Furthermore, the matrix $\mathbf{A}$ being common to each set of samples potentially weakens the resulting MLWE instances. Thus, assuming that $\mathbf{e}_i$ is suitably distributed, it is clear that choosing a γ that is the norm of an element of L compromises security.

We are left to consider the distribution of the error matrix $\mathbf{E}$. In order to understand this, we must discuss the proof of Lemma 5.11 further. Let $\gamma = N_{L/K}(\beta)$, so that the isomorphism mapping $\mathcal{A}$ to $\mathcal{A}' = (L/K, \theta, 1)$ fixes L and sends u to $u\beta$. Following the proof of Theorem 4.3 we see that the γ which are both roots of unity and norm elements are precisely norms of some other root of unity. Hence, β is a root of unity and this isomorphism maps a Gaussian distribution on $\mathcal{A}$ to a Gaussian distribution on $\mathcal{A}'$.

The isomorphism mapping $\mathcal{A}'$ to $M_{d \times d}(K)$ begins with a mapping from $\mathcal{A}'$ to $\text{Hom}_K(L, L)$ that sends $x \in L$ to the multiplication function $f_x(y) = xy$ for $y \in L$ and sends u to θ . Then, it applies the well known isomorphism sending $\text{Hom}_K(L, L)$ to $M_{d \times d}(K)$, which can be defined constructively as follows:

- Fix a K -basis $\{\ell_1, \dots, \ell_d\}$ of L over K .
- Define $f_j : K \rightarrow L$ as $f_j(x) = \ell_j x$, a mapping onto the j coordinate of the basis.

- Let $\pi_j : L \rightarrow K$ denote the projection map onto the ℓ_j sending $\pi_j(\sum_{i=1}^d x_i \ell_i) = \ell_j$.
- Define $\Delta : \text{Hom}_K(L, L) \rightarrow M_{d \times d}(K)$ coordinatewise as $\Delta(\psi)_{i,j} = (\pi_i \circ \psi \circ f_j)(1)$.

Since it permits an arbitrary choice of K -basis, this isomorphism is non-unique. Furthermore, an attacker trying to apply this isomorphism would be able to use their choice of basis and still compute the isomorphism efficiently. We are interested in the image of a Gaussian sample $e \in \Lambda_q$ under this isomorphism, with $e = \sum_{i=0}^{d-1} u^i e_i$, having each e_i sampled independently from a discrete Gaussian over $\mathcal{O}_{Lq}$, being sent to $\psi_e = \sum_{i=0}^{d-1} e'_i \theta^i$. Correspondingly, the i, j coordinate of the matrix $\mathbf{E} = \Delta(e)$ is

$$\pi_i\left(\sum_{k=0}^{d-1} e_k \theta^k(\ell_j)\right).$$

For the j^{th} column of $\mathbf{E}$ (the error vector in the set of d MLWE samples with secret $\mathbf{s}_i$), the error is precisely the ℓ_i coordinate of $\sum_{k=0}^{d-1} e_k \theta^k(\ell_j)$.

Now the distribution of the error in each collection of MLWE samples depends on the properties of the chosen basis. Since the e_k are independent Gaussian samples from L , the j is fixed and θ represents a permutation of the canonical embedding coordinates of L elements. Hence, $\sum_{k=0}^{d-1} e_k \theta^k(\ell_j)$ is an elliptical Gaussian with n blocks of d different parameters. Furthermore, if $\{\ell_1, \dots, \ell_d\}$ is a cyclic basis then, since the distribution of $\|\sigma_L(\ell_k)\|_2$ is independent of k , the projection $\pi_i(\sum_{k=0}^{d-1} e_k \theta^k(\ell_j))$ follows an elliptical Gaussian. In addition, these coordinates are not independent and are potentially highly correlated.

The end result of this exposition is that, depending on the properties of the cyclic bases of L/K and given the choice of γ as a norm element, from a single CLWE instance we can construct d parallel copies of d MLWE instances in dimension n and rank d with correlated error. These correlated instances of the MLWE problem are plausibly substantially easier than the claimed security of the CLWE instance, which is that it is roughly as hard as an MLWE instance in the same dimension nd and rank d . Of course, the error distributions in the underlying MLWE instances are non-standard and we have not presented a concrete attack on them. Instead, we believe this discussion is sufficient to persuade the unconvinced reader that solving the CLWE

problem with norm element γ can be simplified by some parallelization into MLWE instances, and thus we should stick to our specification that γ is a norm.

In the above exposition we restricted ourselves to cases where γ is a norm, but the definition of the non-norm condition also permits γ as valid if and only if γ^i is a non-norm for some $i < d$ that is coprime with d (see [VHLR09]). However, we have previously assumed that the hardness of the CLWE problem was independent of the choice of primitive n^{th} root. In the constructions of Theorem 4.3 γ is an n^{th} root of unity and d divides the prime power n , so if i is coprime with d then i is also coprime with n and so γ^i is a primitive root which defines a cyclic algebra in which the CLWE problem can be parallelized. Thus, we conclude that γ must satisfy the non-norm condition rather than just itself not be a norm. Independently, a recent work [PUTPG⁺19] revisiting m -RLWE observes that the underlying property causing the attacks of [BCV18] on the original instantiations was the presence of zero-divisors in the ambient space. In our case, zero-divisors exist in a cyclic algebra if and only if the non-norm condition is not satisfied, so their argument should preclude not just the γ that are themselves norms but also all γ which fail the non-norm condition.

Chapter 6

Conclusion

Don't adventures ever have an end? I
suppose not. Someone else always has
to carry on the story.

Frodo Baggins, *The Lord of the Rings:
The Fellowship of the Ring* by J.R.R.
Tolkien

6.1 Summary of Thesis Achievements

The primary goal of this thesis is the introduction of the Learning with Errors problem over Cyclic Algebras, CLWE, adding to the family of available LWE assumptions for use in cryptography. To this end, the central pillars of an LWE problem are provided for the cyclic algebra case. First, in order to provide a foundation for the construction the notion of lattices derived from ideals of the natural order of a cyclic algebra are applied in cryptography for the first time. Then, in Chapter 3, the CLWE problem is formally introduced, following which the usual LWE security reductions are established in the CLWE case: namely, the search CLWE problem is reduced to the problem of solving short vector problems on order-ideal lattices, and then the decision CLWE problem is reduced to a variant of the search CLWE problem where the secret

is restricted to a fixed, well constructed subset of its usual space. Under plausible assumptions on this restricted search problem, combining these two reductions gives the necessary security grounding for CLWE based cryptography, which is that samples from the CLWE distribution appear pseudorandom to an onlooker with no knowledge of the secret s . Following this, in Chapter 4, explicit algebras are provided with dimensions and structure appropriate for cryptographic use. Finally, in Chapter 5, the necessary steps are taken to mold the CLWE problem into a practical format for cryptography. Primal and normal form reductions are shown and a sample cryptosystem in this form is provided. Additionally, the complexity of operations in CLWE cryptography is compared to that of RLWE and MLWE based schemes.

6.1.1 Security of CLWE

The CLWE problem is introduced in Chapter 3, and as is typical for algebraic LWE problems it is defined in its dual form. The dual form of the CLWE problem closely mirrors that of the RLWE and MLWE problem: an integer modulus q is chosen, and the natural order takes the place of the discrete ambient subspace from which a is sampled, corresponding to the ring of integers in the RLWE case and the full module over the ring of integers in MLWE. Then, its dual is embedded in a suitable continuous ambient space to set up the decoding problem. As with other LWE instances, the CLWE problem tasks a solver with recovering a secret s from a collection of samples from the CLWE distribution $\Pi_{q,s,\psi}$ which represent noisy ring equations in s (search problem) or distinguishing the distribution $\Pi_{q,s,\psi}$ from uniform (decision problem).

Following the formal definition of CLWE, the exact structured lattice problems that it is reduced to are introduced. Once again, the lattice problems presented are the cyclic algebra equivalent of those used in RLWE. An order-ideal lattice is defined as the lattice in H_L^d derived from applying the embedding $\sigma_{\mathcal{A}}$, the canonical embedding of L applied independently on each of the u coordinates of $\mathcal{A}$, to a two-sided ideal $\mathcal{I}$ of the natural order Λ . On these lattices the standard collection of lattice problems are defined: first, the (approximate) order-ideal Shortest Vector Problem is the problem of finding a (nearly) shortest element $a \in \mathcal{I}$ in the prescribed norm, which is typically the ℓ_2 norm on the image of $\sigma_{\mathcal{A}}$. The order-ideal Shortest Independent

Vectors Problem is the problem of finding N linearly independent elements $a_1, \dots, a_N$ whose largest element has norm (almost) as short as possible. Since the embedding $\sigma_{\mathcal{A}}$ considers the structured lattices as integer lattices, the output elements are required to be linearly independent over $\mathbb{Z}$. Finally, the order-ideal Bounded Distance Decoding problem is the problem of decoding an element of $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ to the nearest element of a two-sided ideal $\mathcal{I}$ of the natural order given the guarantee that the distance is bounded above by some positive δ . That is, given $y = x + e$ for $x \in \mathcal{I}$ and $e \in \bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ satisfying $\|\sigma_{\mathcal{A}}(e)\|_2 \leq \delta$, the order-ideal BDD problem is to find x .

Next, justifications for certain design choices made in the definition of the CLWE problem are discussed in some detail. The order of operations is chosen as $a \cdot s$ rather than $s \cdot a$ to continue the narrative of algebraic LWE problems as structured instances of the noisy matrix-vector equation $\mathbf{b} = \mathbf{A}s + \mathbf{e}$ over $\mathbb{Z}_q$. Choosing the order of multiplication to have a on the left structures $\mathbf{A}$ by setting $\mathbf{A} = \text{rot}(\phi(a))$, where the rotational operation used in RLWE is applied coordinatewise on the elements of the matrix $\phi(a)$ as in the MLWE case. The natural order is chosen, rather than some maximal order, to continue the comparison to MLWE in which the module used is typically the full module $\mathcal{O}_K^d$. Additionally, the use of the natural order has a substantial advantage in terms of conceptual and representational simplicity. The pair of number fields is chosen so that the dimension of K scales asymptotically and L is a low degree extension for two reasons: in part to further the analogy with MLWE by keeping the equivalent of module rank as a small constant, and also to permit the search-decision reduction which runs in polynomial time only in $n = [K : \mathbb{Q}]$ but not $d = [L : K]$. The restriction of γ to non-norm elements, so that the algebra $\mathcal{A}$ is a division algebra, is made to preemptively prevent parallelizing attacks in the style of that of [BCV18] on the multivariate RLWE problem. This choice is expanded upon later in Section 5.4 once some necessary technical steps have been taken, where weak instances of the CLWE problem are transformed into parallel instances of the MLWE problem over K with a non-standard error distribution.

Following the discussion, the remainder of the section is spent showing that the usual security properties for LWE problems hold in the case of CLWE. First, the reduction from the search CLWE problem to the order-ideal SVP problem is presented, following the usual outline from

[Reg05] for the plain LWE problem. The process begins by reducing the Discrete Gaussian Sampling problem over order-ideal lattices to order-ideal SVP using techniques from [LS15] for module lattices. Then, the standard ‘bootstrapping’ step reduces the BDD problem over an order-ideal to the Discrete Gaussian Sampling problem. The last step of the reduction is to transform an order-ideal BDD sample $y = x + e'$ to a CLWE instance $(a, a \cdot s + e)$, where s corresponds to x , following which a search CLWE oracle can be used to solve the BDD instance. This step requires the most technical work to adapt to the CLWE case, including a collection of lemmas (Section 3.4) allowing order-ideals to be ‘cleared out’ just as ring ideals are in the corresponding RLWE security reduction in [LPR10].

After the hardness of the search CLWE problem has been derived from the assumed difficulty of solving the order-ideal SVP problem, the next step is to show the corresponding hardness of the decision problem. The search-decision reduction is a close adaptation of that of [LPR10]. Here, the main technical innovation is to apply the CRT style decomposition of [OAS12] given in Lemma 2.34 to the special case of the ideal generated by the modulus q in $\mathcal{O}_K$, the use of which specifies that K should be a cyclotomic field within which q splits completely. Using this CRT decomposition requires an additional assumption that the hardness of the CLWE problem is independent of the choice of primitive root of unity γ , since the K -automorphisms permuting the ideals in the decomposition of q are extended to $\mathcal{A}$ in a manner that permutes the primitive roots. Otherwise, the techniques in the reduction are primarily analogs of a combination of the ring and module search-decision reductions. The main exception is that when it comes to ‘guessing’ the secret in one slot of the CRT decomposition the structure of multiplication in a cyclic algebra means that the reduction complexity is no longer polynomial in d . Additionally, this step restricts the search problem to one in which $s \in G \subset \Lambda_q^\vee$ holds for any well chosen set G whose image under each coordinate of Lemma 2.34 satisfies that the pairwise difference of any two elements is invertible. It is assumed that the resulting search problem is still hard, since the reduction is valid for any G , and thus the decision CLWE problem is as hard as the search variant for the hardest choice of G .

To conclude Chapter 3, the difficulty of adapting the direct ideal-BDD to decision RLWE reduction from [PRSD17] to the case of CLWE is briefly discussed. Unlike in the ring and

module cases, the coordinates of a cyclic algebra can not be sufficiently separated to apply their techniques, which essentially ‘guess’ the error e coordinatewise and use a hill climbing algorithm independently in each coordinate to find the correct value. Unfortunately, this requires that perturbing one coordinate of the BDD instance $y = x + e$ only effects one coordinate of its image under the transformation into a CLWE instance. This is true of RLWE by considering coordinates of the canonical embedding, but there does not appear to be a corresponding orthogonal representation of multiplication in the cyclic algebra setting.

6.1.2 Constructing Appropriate Algebras

The next section is spent constructing suitable algebras to instantiate the CLWE problem in. First, the restrictions placed on L , K , and γ induced by the definition of CLWE and the security reductions of Chapter 3 are made explicit: namely, γ must be an integral non-norm element to maintain a well-defined cyclic division algebra with closed natural order and each L -embedding $\sigma_i(\gamma)$ should have absolute value 1 to maintain sub-multiplicativity of the norm. This condition is simplified to restricting γ to the roots of unity in K that are non-norm elements of the extension L/K . The degree n of K should scale asymptotically since the security reductions are efficient in n ; correspondingly, L should be a low degree extension of K . Finally, K should be a cyclotomic field within which the modulus q splits completely.

Given these restrictions, the techniques of Chapter 4 begin with the existing construction of Theorem 4.1, taken from [LMM08]. This theorem constructs a cyclic division algebra with center $K = \mathbb{Q}(\zeta_m)$ and $[L : K] = m$, where $m = p^a$ is a prime power. Furthermore, any primitive m^{th} root of unity $\gamma = \zeta_m$ satisfies the non-norm condition in the extension L/K . These algebras are not immediately suitable for the CLWE problem because the degree d is necessarily larger than $n := \phi(m)$. However, Theorem 4.3 demonstrates how to use techniques from Galois theory to create a subextension $L/M/K$ of any degree dividing m in which γ is still a non-norm element. Then, Section 4.3 provides example parameter sets where the algebras constructed by Theorem 4.3 achieve suitable dimension for cryptographic use that satisfy the restrictions that K is a large cyclotomic field and $[L : K]$ is small. The target

dimensions are chosen by taking motivation from parameters of algebraic LWE based NIST second round candidates such as NewHope and Kyber. Section 4.2.2 shows an alternative way to use Theorem 4.1 to construct low degree extensions of a large center: pick d as a small prime power and construct a small field extension L/K using Theorem 4.1, then increase dimension by composing L and K with a large cyclotomic field M that is linearly disjoint from L . This method is used to construct appropriately sized fields for cryptographic use in which K is a cyclotomic field but no longer a prime power.

After this, the section is rounded off by considering cases where q splits completely into a product of prime ideals in both L and K , a property used in Section 5.3 to support fast multiplication. To identify these cases the proof of Theorem 4.1 is examined: all fields constructed using this technique are subfields of $L' = K(\zeta_s)$ for some large prime number s , so q splits completely in any subfield of L if it splits completely in L' . This method constructs splitting q which are impractical due to their size, so more streamlined techniques are considered for specific cases. For example, in the quadratic case where L is the unique degree 2 extension of K that is a subfield of L' there are standard methods to compute splitting in L .

6.1.3 CLWE in Cryptography

In order to apply the CLWE assumption in an appropriate shape for use in cryptography, the samples must be discretized and the secret should be small. Thus, the beginning of Chapter 5 is spent constructing the primal form of CLWE, discretizing the samples of the CLWE distribution to the ambient space $\Lambda_q \times \Lambda_q$ by using a simple adaptation of the technique for RLWE of [LPR13] to cyclic algebras. Following this, the normal form is shown to be secure, so that the secret s can be sampled from the now discretized error distribution χ . The algorithm is the same in spirit as in the case of classical LWE in [ACPS09], with the main technical step required being an observation that a non-negligible proportion of the elements of Λ_q are invertible, which follows by an appeal to Lemma 2.34 and then Wedderburn's Theorem.

After the primal and normal forms have been established, the cyclic algebra variant of the classical LWE cryptosystem, described in its ring form in Section 2.7, is defined. The security

of the system is shown to derive from the assumed hardness of the CLWE problem, and the correctness is shown to follow from the usual properties of the error distribution used in the underlying CLWE instances. The main novel technique required to instantiate this scheme in a cyclic algebra is Lemma 5.6, which shows that applying the matrix transpose operation to $\phi(a)$ does not compromise the integrity of CLWE samples.

To supplement the CLWE based cryptosystem, fast algorithms are provided to compute the requisite multiplication operations. These algorithms are a combination of the CRT style decomposition of Lemma 2.34 and existing algorithms for multiplication in a skew polynomial ring from [CLB17]. First, it is shown under plausible restrictions on L , K , and q that the rings $\mathcal{R}_i$ constructed by Lemma 2.34 satisfy the conditions required by the algorithms of [CLB17]. Then, the complexity of the decomposition into the rings $\mathcal{R}_i$ is computed using the CRT algorithm of [LPR13] on the cyclotomic field K and handling the non-cyclotomic part L by constructing suitable bases whose image under the CRT algorithm can be pre-computed. This technique is exhibited in the case where q splits completely in L , where it is particularly efficient, since given a completely split q it is straightforward to find an orthonormal $\mathcal{O}_{Kq}$ basis of $\mathcal{O}_{Lq}$ using the Chinese Remainder Theorem on the decomposition of q in L and K respectively. Finally, the technique is generalized to the case of arbitrary bases and splitting of q at a small increase in complexity.

6.2 Discussion and Future Work

This thesis lays the groundwork for the use of the Cyclic Learning With Errors problem, and more generally cyclic algebras, in cryptography. Now that they have been formally introduced, cyclic algebras exhibit substantial novel structure within lattice based cryptography, and discovering use cases for these previously unseen features represents an exciting area of future research.

6.2.1 Full Hardness of the Decision Problem

From a theoretical standpoint, the most pressing question to be solved about CLWE is whether or not the search and decision problem are polynomial time equivalent, or instead if the hardness of the decision variant can be based directly on hard lattice problems via some other technique. In this work, the hardness of the decision problem for arbitrary secret is shown to derive from the assumed hardness of a variant of the search problem where the secret is restricted to lie in any so-called pairwise difference set G . Although this substantially lowers the size of the secret space, the resulting secret space is still far too large to exhaustively search. Furthermore, the decision problem is as hard as the search problem for the hardest choice of decision set G , precluding particularly easy cases. Nonetheless, this does not establish the formal hardness of the decision CLWE problem based on the lattice problems of Section 3.2. The reduction fails to permit an arbitrary secret since the decomposition into matrix rings of Lemma 2.34 results in a problem that can not be ‘guessed’ effectively, since the oracle does not necessarily accept inputs as valid when the guess is wrong.

Another method of establishing the hardness of decision RLWE that is not shown for CLWE in this work is a direct-to-decision reduction, which more generally represents a security proof for the decision problem that holds for wider classes of cyclic division algebras than those of Theorem 3.20. The direct-to-decision reduction of [PRSD17] is the only security reduction for RLWE which establishes the hardness of the decision problem without enforcing that K is a cyclotomic field within which q splits completely, as in the search-decision reduction of [LPR10] and the presented analog for CLWE. Dropping this restriction, and hence widening the possible choices of cyclic algebras supporting the hardness of the decision problem, would provide larger design space for CLWE based cryptography. However, there remain substantial technical hurdles to overcome. The difficulty of adapting [PRSD17] was discussed in Section 3.6.6, although it remains plausible that the core idea, that of hill climbing various coordinates of the BDD error separately, can somehow be adapted to the CLWE case.

It is plausible that the inability to show the hardness of decision CLWE as a lattice problem via either of the known methods for RLWE is a result of a more causal issue with CLWE, that

the decision problem is somehow tangibly easier than the search variant. The link between the issues in applying the two methods is that the CLWE problem is substantially less easy to separate algebraically into coordinatewise representations than the existing variants, rather than some aspect of the non-commutative structure. Perhaps this non-linearity opens up a weakness that an attacker could exploit, but the work of [PP19] establishes the hardness of the decision variant of the Middle-Product LWE problem by reducing it to RLWE, despite the fact that the Middle-Product LWE problem also displays a structure that does not parallelize well. In general, there is no obvious reason that a problem that is hard to separate algebraically should be easier than one that splits easily. Thus, it remains open, and indeed unclear, whether or not an entirely novel technique can establish formal hardness of the decision problem. One method that is not explored in this work is to try to utilize the new machinery of [PP19] that reduces several other structured LWE problems to the RLWE problem over number fields of appropriate dimension for the case of decision CLWE, since a valid reduction to an RLWE problem would certainly imply security based on structured lattice assumptions.

6.2.2 Fleshing Out CLWE

Perhaps the most natural line of future research is to study the construction choices and basic applications of the CLWE problem further. The most obvious direction in which to take this is to design an explicit, practical cryptosystem in the style of Section 5.2. That is, to select concrete parameters targeting specific security levels, such as the security levels specified by the NIST process, in the style of NewHope for RLWE and Kyber for MLWE. The formalization of the scheme requires a choice of K , L , γ , representations of elements, error distributions, error correcting codes or lack thereof, as well as considerations of practical security measures such as constant time algorithms and sources of randomness. Additionally, the multiplication algorithms of Section 5.3 must be converted from theory to practice, in the same way that the NTT-based multiplications of [ADPS16] and [BDK⁺18] are practical realizations of the CRT decomposition and coordinatewise multiplication technique used in ring and module LWE schemes. Only once a parameterized scheme is suggested can the practical efficiency of CLWE

be compared to that of existing LWE constructions.

An interesting ‘gap’ in the security proofs of Chapter 3 is that the security of the CLWE problem is based on the hardness of lattice problems derived from embeddings of two-sided ideals of Λ , even though the problem somewhat naturally corresponds to that of decoding on a right ideal of Λ . In terms of explicitly solving the CLWE problem, presumably one would take into account the fact that $a \cdot s$ is certainly in the right span of a , but not necessarily the left. Defining $\langle x \rangle_R$ as the smallest right ideal of Λ containing x , it should be clear that b is relatively close to one element of $\langle a \rangle_R$, the right ideal generated by the public element a . Explicitly, $b := a \cdot s + e$ is close to the element $a \cdot s$, an element of the right ideal generated by a , as long as e is small enough. The difficulty in deriving security from right ideals stems from a lack of analogs of both their decomposition modulo q as in Lemma 2.34 for two-sided ideals and a ‘clearing out’ lemma for right ideals as in Section 3.4. Alternatively, with many samples one could consider solving CLWE as a ‘hidden left ideal’ problem, where the samples are close to uniformly random samples from the left ideal $\langle s \rangle_L$. However, whether or not this is sufficiently distinct from the ring case, where one replaces the left ideal generated by s with the principal ideal $\langle s \rangle$, to provide additional attack surface is unclear. These observations do not invalidate the security of the CLWE problem, but some notion of security derived from right ideals or left ideals would be a strictly stronger notion than from two-sided ideals, since the set of two-sided ideals is the intersection of that of right ideals and left ideals.

One interesting observation on this topic is that the use of the ‘clearing out’ Lemma 3.7 can be restricted to the case where two of the ideals are naturally two-sided, since in practice they correspond to $\langle q \rangle$ and $\Lambda^\vee$. The issue with presenting a clearing out lemma for one-sided ideals is the complexities of one-sided quotients, but perhaps it is possible to clear out a single one-sided ideal $\mathcal{I}$, corresponding to the lattice the BDD problem occurs in, from a pair of two-sided ideals without having to resort to complicated algebra involving the one-sided ideals themselves. Once one is able to clear out ideals, it should be easy to show a transformation from a one-sided BDD instance to the CLWE problem. Following that, one-sided ideals are still modules and so the remainder of the reduction to short vector problems over the appropriate lattices should follow in exactly the same manner as in the two-sided case.

In a different direction, the CLWE problem as presented is instantiated in the natural order of a cyclic algebra. However, it may be interesting to try to construct the CLWE problem over a maximal order instead. Since a maximal order of a cyclic algebra causally has a higher density than the natural order, intuitively one would expect that the corresponding CLWE problem is harder due to the increased difficulty of the underlying decoding problem. Adapting the search security reduction should be straightforward, because the two-sided ideals of the maximal order form a free abelian group and hence lemmas corresponding to those of Section 3.4 should follow in much the same way as in the ring case. However, the search-decision reduction relies on the CRT-like decomposition in the natural order, whereas applying the Chinese Remainder Theorem constructively in the case of prime ideals of the maximal order does not have the same existing precedent. Perhaps a family of cyclic division algebras with a closed form family of two-sided ideals, or more simply a closed form factorization of $\langle q \rangle$, can be found that permits a search-decision reduction using familiar techniques. For example, one could consider a maximal order which contains the natural order as a subring and extend the CRT style decomposition in some manner. Alternatively, an adaptation of the direct-to-decision reduction seems as likely to hold in the case of a maximal order as it does in the natural order and would bypass the issues based on the CRT.

Finally, one might consider if some of the restrictions placed on γ can be dropped. Following the exposition of Section 5.4 it is unlikely that one can safely drop the requirement that γ is a non-norm element. Additionally, recall that γ must be integral and satisfy that $|\sigma_i(\gamma)| = 1$ for each L -embedding σ_i . There are two ways in which one might be able to drop the integrality of γ , which is required to make the natural order well-defined. The first is discussed above; any cyclic division algebra contains at least one maximal order, so one could consider setting γ to be some non-integral element and instantiating CLWE in a maximal order so that the loss of the natural order is not so dramatic.

Alternatively, one could consider using the relatively new construction of [PP19]. In [PP19], the authors generalize algebraic forms of the LWE problem and introduce the notion of a coefficient ring. Their generalization, which they denote $\mathcal{L}$ -LWE, samples s from (the dual of) some lattice $\mathcal{L}$ in the ambient space, a vector space over K , over which the LWE operation $a \cdot s$

can be considered as a K -bilinear form. A semantic adaptation of their Definition 3.7 permits the CLWE problem as a valid bilinear form over its center K but not the maximal subfield L . The novel idea is that a is not necessarily sampled from the same lattice as s but instead from its coefficient ring $\mathcal{O}^{\mathcal{L}} := \{x \in K : x\mathcal{L} \subset \mathcal{L}\}$. Replacing K in this definition with $\mathcal{A}$, one can consider sampling s from some lattice $\mathcal{L}$ in $\mathcal{A}$ which is not closed under multiplication and then sampling a from its coefficient ring rather than $\mathcal{L}$ itself. The dropping of the restriction that the secret should be sampled from a lattice which is closed under multiplication means one may be able to replace γ with some element which is not integral but still satisfies $|\sigma_i(\gamma)| = 1$ for each i whilst still sampling s from $\bigoplus_i u^i \mathcal{O}_{L_Q}$. For example, the choices of γ from [ESK07] give suitable γ which are fractional elements of K in order to achieve so-called perfect codes, and algebras with their choices of γ satisfy certain desirable properties for coding which may also be of use in cryptography.

Of course, substantially more design space would exist were γ not restricted to have $|\sigma_i(\gamma)| = 1$, enforced to maintain that the ℓ_2 norm induced by the $\sigma_{\mathcal{A}}$ embedding is a well-defined sub-multiplicative norm. One could plausibly drop this condition on γ by changing the norm to a weighted ℓ_2 norm on $\sigma_{\mathcal{A}}$. However, this weighting has to take into account two properties: higher u^i coordinates need weighting to take into account that they will be acted on by γ more often during a multiplication step. More intricately, one must weight coordinates of each of the L -embeddings σ_i relative to the value of γ in that coordinate, which has to be done blockwise in each of the extensions of the K -embeddings. This can only be done in a well-defined manner because $\gamma \in K$ holds and thus is constant on each of the blocks corresponding to all L extensions of one K embedding.

For example, the norm of an element $a \in \mathcal{A}$ could be defined by taking some fixed, arbitrary, choice of L -extension $\alpha = (\alpha_1, \dots, \alpha_n)$ of the K -canonical embedding and then taking the ℓ_2 norm of $\alpha(\phi(a))$, where α is applied coordinatewise. This lines up with the ℓ_2 norm used throughout this work when each $|\sigma_i(\gamma)| = 1$ but considers the effect of γ once this condition is dropped. One can even sample an elliptical Gaussian $D_{\mathbf{r}}$ over $\bigoplus_{i=0}^{d-1} u^i L_{\mathbb{R}}$ in this norm by

taking an elliptical Gaussian $D_{\mathbf{r}'}$ over $L_{\mathbb{R}}^d$ in the ℓ_2 norm for

$$r'_{i_j} = \frac{r_{i_j}}{(d-i) + \sum_{k=1}^i |\sigma_{\Delta^{d-i+k}(j)}(\gamma)|},$$

where Δ denotes the permutation of coordinates of the canonical embedding induced by θ . This takes into account the weight of γ , but even if the same security properties can be shown in this norm it is likely that the extra flexibility in choice of γ is not worth the technical complications and the increasingly non-standard lattice problems which would be induced.

6.2.3 RLWE Attributes in CLWE

The next obvious place to look for future work is to see which properties of module and ring based cryptography that have not been discussed in this work can be adapted to cyclic algebras. Although the classical LWE problem is believed to be truly hard in a post quantum world, the algebraic structure of RLWE has been subject to substantial cryptanalysis in recent years that does not effect the security of classic or module LWE. In this work the hardness of the CLWE problem is conjectured to lie between that of the MLWE problem and the RLWE problem, but quantifying this statement remains an interesting problem: is the concrete security closer to that of ring or module LWE, and how hard are order-ideal lattice problems compared to their more established counterparts?

The majority of the RLWE cryptanalysis follows one of two avenues, relating to properties of the underlying number field K . First, there is a sequence of quantum attacks (see e.g. [EHKS14], [CGS14], [BS15], [CDW17]) taking advantage of the structure of the unit group of the ring of integers $\mathcal{O}_K$. The unit groups of cyclic algebras, and more generally of non-commutative division algebras, are much harder objects to work with than their number field equivalents. Primarily, since they are non-commutative, common techniques based on the generalizations of the Abelian Hidden Subgroup Problem applied to the unit group of the ring of integers can not be brought to bear on them. Secondly, there do not exist efficient algorithms that compute the unit group of Λ , nor is there a known low-index subgroup forming an analog of the cyclotomic

units used in [CDW17] for the ring case. Finally, although in general there will be more units in the cyclic algebra of the same dimension than in a number field, the density of such units is far lower. For example, if $a, b \in \mathcal{O}_L$ are units in L then each $u^i a$ is a unit in any cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ in which γ is an $\mathcal{O}_K$ -unit, but the combination $u^i a + u^j b$ need not be. Hence, $\mathcal{A}$ has at least d times as many units as L but has exponentially more elements. Understanding how to adapt these unit group based techniques, if indeed it is even possible, would go some way to clarifying the gap between ring and cyclic LWE security.

The other avenue of ring cryptanalysis is based on utilizing the ring automorphisms of K . Although this is often combined with methods involving the unit group, the underlying methods of this technique may be implementable in cyclic algebras and thus should be considered separately. Extending the Galois automorphisms of K to $\mathcal{A}$ as in Theorem 3.20 corresponds to an action on the Brauer group of $\mathcal{A}$. Brauer Groups and Morita equivalence classes of central simple algebras, and in particular cyclic algebras, are deeply structured and well studied objects, so perhaps tools used in this study can be applied in place of automorphisms of $\mathcal{A}$ to attack the algebraic structure.

Alternatively, one might try to attack the structured lattice problems directly. This seems the least likely method of cryptanalyzing CLWE based schemes; the order-ideal lattices are less structured than ring ideal lattices and thus the corresponding problems should be harder. Furthermore, the aforementioned attacks on RLWE are typically on related lattice problems rather than direct attacks on ideal SVP and so it seems unlikely that direct attacks on order-ideal SVP would be found given that no substantial progress has been made on the most likely easier ring variant. Finally, as in RLWE, solving the underlying lattice problems does not necessarily solve the CLWE problem, since the reductions in this work are only one-way.

Finally, the majority of ring and module LWE based cryptosystems have concrete multiplication algorithms based on the NTT, essentially a practical application of the CRT decomposition technique of [LPR13]. The fastest multiplication algorithm of Section 5.3 applies only for the family of algebras where q splits completely in both L and K , but can be extended to the general case at only a small cost. It is unclear whether this multiplication algorithm

is the best possible, or if a specialized CLWE algorithm would achieve better results. The algorithm of Section 5.3 essentially just combines the CRT decomposition of [LPR13] and the skew polynomial multiplication algorithms of [CLB17]. Whilst both of these algorithms apply to the case of CLWE, it may be possible to build bespoke algorithms for cyclic algebras which outperform this technique.

6.2.4 More Cryptographic Functionality

The use of cyclic algebras in cryptography is the principal novel contribution of this thesis. The main motivation for restricting the use of cyclic algebras to CLWE is that LWE assumptions support well understood security reductions to lattice vector problems. By defining an LWE assumption, the hardness of cryptography instantiated in a cyclic algebra can be somewhat quantified before the years of scrutiny typically offered to cryptosystems are applied. However, there is no a priori reason why other lattice based schemes could not be adapted to the case of cyclic algebras. The majority of lattice cryptosystems not based on LWE problems are in the style of NTRU, which takes place over polynomial rings. Taking the example of Falcon [FHK⁺17], an NTRU based signature scheme over cyclotomic rings in the second round of the NIST process, it is likely that cyclotomic ring based NTRU schemes can be adapted to work over cyclic algebras in the same way that cyclotomic ring LWE was adapted to CLWE. However, the security of NTRU schemes is more enigmatic than that of LWE schemes and thus cryptographers would likely require some convincing that instantiating NTRU in cyclic algebras does not introduce additional tools for cryptanalysts to take advantage of.

Moving on from the basic construction of public key schemes, one could investigate whether CLWE enables some elements of so called ‘crypto dream-land’ to become reality. Cryptographers have too many wishes to list exhaustively here, but a subject of particular focus is constructing Functional Encryption, and with it Indistinguishability Obfuscation (see for example [BV18]). Fledgling techniques for this seem to be based on the plain LWE assumption, but perhaps the novel non-commutative algebraic structure may be of some use. More plausibly, one could return to the now defunct construction of multilinear maps from [GGH13],

broken in [CJL16] among other papers. The techniques for constructing multilinear maps use the algebraic structure of RLWE and consequentially open themselves up to attacks using the number field structure. This makes them a promising candidate for considering novel properties of cyclic algebras in cryptography, since as seen in this work there are times when the cyclic algebra and number field structures feel remarkably similar and times when they feel very far apart. Thus, an optimistic cryptographer might hope that they are similar enough that CLWE can be used to construct multilinear maps yet distinct enough that the attacks do not lift to the non-commutative case.

6.3 Concluding Remarks

In this thesis the CLWE problem has been rigorously introduced to the world of cryptography. The requisite security properties are established, and through explicit constructions and multiplication algorithms initial moves towards a practical CLWE based cryptosystem have been made. However, since the properties shown in this thesis represent the bare bones requirements, this is just the beginning of the story of cyclic algebras in cryptography. What more can be done with non-commutative lattice cryptography? Building a practical cryptosystem supporting concrete parameters and implementations in the style of the NIST candidates is the logical first step, but beyond that there seems no clear road map determining what properties of cryptographic constructions based on cyclic algebras might arise in the future; this thesis leaves behind a new, large, and exciting research area to be explored.

Bibliography

- [ABD16] Martin Albrecht, Shi Bai, and Léo Ducas. A subfield lattice attack on over-stretched NTRU assumptions. In *Annual International Cryptology Conference*, pages 153–178. Springer, 2016.
- [ABD⁺18] Erdem Alkim, Joppe Bos, Léo Ducas, Patrick Longa, Ilya Mironov, Michael Maehrige, Valeria Nikolaenko, Chris Peikert, Ananth Raghunathan, and Douglas Stebila. FrodoKEM. <https://frodokem.org/>, 2018.
- [ABD⁺19] Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Kyber algorithm specifications and supporting documentation (version 2.0). <https://pq-crystals.org/kyber/data/kyber-specification-round2.pdf>, 2019.
- [ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Advances in Cryptology-CRYPTO 2009*, pages 595–618. Springer, 2009.
- [AD97] Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 284–293. ACM, 1997.
- [AD17] Martin R Albrecht and Amit Deo. Large modulus ring-LWE $\geq$ module-LWE. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 267–296. Springer, 2017.

- [ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange—A New Hope. In *25th USENIX security symposium*, pages 327–343. USENIX Security, 2016.
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 99–108. ACM, 1996.
- [BBF⁺18] Hayo Baan, Sauvik Bhattacharya, Scott Fluhrer, Oscar Garcia-Morchon, Thijs Laarhoven, Rachel Player, Ronald Rietman, Markku-Juhani Saarinen, Ludo Tolhuizen, Jose Luis Torre Arce, and Zhenfei Zhang. Round5. <https://round5.org/>, 2018.
- [BBPS18] Madalina Bolboceanu, Zvika Brakerski, Renen Perlman, and Devika Sharma. Order-LWE and the hardness of Ring-LWE with entropic secrets. Cryptology ePrint Archive, Report 2018/494, 2018. <https://eprint.iacr.org/2018/494>.
- [BCD⁺16] Joppe Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! Practical, quantum-secure key exchange from LWE. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1006–1018. ACM, 2016.
- [BCLvV17] Daniel Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Christine van Vredendaal. NTRU Prime: reducing attack surface at low cost. In *International Conference on Selected Areas in Cryptography*, pages 235–260. Springer, 2017.
- [BCV18] Carl Bootland, Wouter Castryck, and Frederik Vercauteren. On the security of the Multivariate Ring Learning With Errors problem. Cryptology ePrint Archive, Report 2018/966, 2018. <https://eprint.iacr.org/2018/966>.
- [BDK⁺18] Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-

- Kyber: a CCA-secure module-lattice-based KEM. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 353–367. IEEE, 2018.
- [Ber14] Daniel Bernstein. A subfield-logarithm attack against ideal lattices. <http://blog.cr.yp.to/20140213-ideal.html>, 2014.
- [BJMM12] Anja Becker, Antoine Joux, Alexander May, and Alexander Meurer. Decoding random binary linear codes in $2^{n/20}$: how $1 + 1 = 0$ improves information set decoding. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 520–536. Springer, 2012.
- [Bra16] Matt Braithwaite. Experimenting with post-quantum cryptography. <https://security.googleblog.com/2016/07/experimenting-with-post-quantum.html>, 2016.
- [BS15] Jean-François Biasse and Fang Song. On the quantum attacks against schemes relying on the hardness of finding a short generator of an ideal in $\mathbb{Q}(\zeta_{p_n})$. <http://cacr.uwaterloo.ca/techreports/2015/cacr2015-12.pdf>, 2015.
- [BS16] Jean-François Biasse and Fang Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 893–902. SIAM, 2016.
- [BV18] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. *Journal of the ACM (JACM)*, 65(6):39, 2018.
- [Cas17] Davide Castelvecchi. IBM’s quantum cloud computer goes commercial. *Nature News*, 543(7644):159, 2017.
- [CDH⁺18] Cong Chen, Oussama Danba, Jeffrey Hoffstein, Andreas Hülsing, Joost Rijnveld, John Schank, Peter Schwabe, William Whyte, and Zhenfei Zhang. NTRU. <https://ntru.org/>, 2018.

- [CDPR16] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 559–585. Springer, 2016.
- [CDW17] Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short Stickelberger class relations and application to Ideal-SVP. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 324–348. Springer, 2017.
- [CGS14] Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: A cautionary tale. In *ETSI 2nd quantum-safe crypto workshop*, pages 1–9, 2014.
- [CJL16] Jung Hee Cheon, Jinhyuck Jeong, and Changmin Lee. An algorithm for NTRU problems and cryptanalysis of the GGH multilinear map without a low-level encoding of zero. *LMS Journal of Computation and Mathematics*, 19(A):255–266, 2016.
- [CLB17] Xavier Caruso and Jérémy Le Borgne. Fast multiplication for skew polynomials. In *Proceedings of the 2017 ACM on International Symposium on Symbolic and Algebraic Computation*, pages 77–84. ACM, 2017.
- [CP16] Eric Crockett and Chris Peikert. Challenges for Ring-LWE. Cryptology ePrint Archive, Report 2016/782, 2016. <https://eprint.iacr.org/2016/782.pdf>.
- [CS97] Don Coppersmith and Adi Shamir. Lattice attacks on NTRU. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 52–61. Springer, 1997.
- [DKRV18] Jan-Pieter D’Anvers, Angshuman Karmakar, Sujoy Sinha Roy, and Frederik Vercauteren. Saber: Module-LWR based key exchange, CPA-secure encryption and CCA-secure KEM. In *International Conference on Cryptology in Africa*, pages 282–305. Springer, 2018.

- [DR99] Joan Daemen and Vincent Rijmen. AES proposal: Rijndael. <https://csrc.nist.gov/csrc/media/projects/cryptographic-standards-and-guidelines/documents/aes-development/rijndael-ammended.pdf>, 1999.
- [EHKS14] Kirsten Eisenträger, Sean Hallgren, Alexei Kitaev, and Fang Song. A quantum algorithm for computing the unit group of an arbitrary degree number field. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 293–302. ACM, 2014.
- [ESK07] Petros Elia, BA Sethuraman, and P Vijay Kumar. Perfect space–time codes for any number of antennas. *IEEE Transactions on Information Theory*, 53(11):3853–3868, 2007.
- [FHK⁺17] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregory Seiler, William Whyte, and Zhenfei Zhang. Falcon. <https://falcon-sign.info/>, 2017.
- [FO99] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In *Annual International Cryptology Conference*, pages 537–554. Springer, 1999.
- [Gab85] Ernest Mukhamedovich Gabidulin. Theory of codes with maximum rank distance. *Problemy Peredachi Informatsii*, 21(1):3–16, 1985.
- [GGH13] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 1–17. Springer, 2013.
- [Gib17] Elizabeth Gibney. D-wave upgrade: How scientists are using the world’s most controversial quantum computer. *Nature News*, 541(7638):447, 2017.
- [Ham17] Mike Hamburg. Post-quantum cryptography proposal: Threebears. <https://sourceforge.net/projects/threebears/>, 2017.

- [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H Silverman. NTRU: A ring-based public key cryptosystem. In *International Algorithmic Number Theory Symposium*, pages 267–288. Springer, 1998.
- [Joz01] Richard Jozsa. Quantum factoring, discrete logarithms, and the hidden subgroup problem. *Computing in Science & Engineering*, 3(2):34–43, 2001.
- [KF17] Paul Kirchner and Pierre-Alain Fouque. Revisiting lattice attacks on over-stretched NTRU parameters. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 3–26. Springer, 2017.
- [LLZ⁺18] Xianhui Lu, Yamin Liu, Zhenfei Zhang, Dingding Jia, Haiyang Xue, Jingnan He, Bao Li, Kunpeng Wang, Zhe Liu, and Hao Yang. LAC: Practical Ring-LWE based public-key encryption with byte-level modulus, 2018. <https://eprint.iacr.org/2018/1009.pdf>.
- [LMM08] J. Lahtonen, N. Markin, and G. McGuire. Construction of multiblock space-time codes from division algebras with roots of unity as non-norm elements. *IEEE Transactions on Information Theory*, 54(11):5231–5235, 2008.
- [LMvdP15] Thijs Laarhoven, Michele Mosca, and Joop van de Pol. Finding shortest lattice vectors faster using quantum search. *Designs, Codes and Cryptography*, 77(2-3):375–400, 2015.
- [LN16] Dirk Liebholt and Gabriele Nebe. Automorphism groups of Gabidulin-like codes. *Archiv der Mathematik*, 107(4):355–366, 2016.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and Learning With Errors over rings. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 1–23. Springer, 2010.
- [LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for Ring-LWE cryptography. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 35–54. Springer, 2013.

- [LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, 2015.
- [LS19] Vadim Lyubashevsky and Gregor Seiler. NTTRU: truly fast NTRU using NTT. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2019(3):180–201, 2019.
- [LvdPdW12] Thijs Laarhoven, Joop van de Pol, and Benne de Weger. Solving hard lattice problems and the security of lattice-based cryptosystems. *Cryptology ePrint Archive*, 2012. <https://eprint.iacr.org/2012/533.pdf>.
- [LVL18] Laura Luzzi, Roope Vehkalahti, and Cong Ling. Almost universal codes for MIMO wiretap channels. *IEEE Transactions on Information Theory*, 64(11):7218–7241, 2018.
- [MAB⁺19] Carlos Aguilar Melchor, Nicolas Aragon, Slim Bettaieb, Loïc Bidoux, Olivier Blazy, Alain Couvreur, Jean-Christophe Deneuville, Philippe Gaborit, Adrien Hauteville, and Gilles Zémor. RQC. <https://pqc-rqc.org/>, 2019.
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.
- [NCS⁺17] Florian Neukart, Gabriele Compostella, Christian Seidel, David Von Dollen, Sheir Yarkoni, and Bob Parney. Traffic flow optimization using a quantum annealer. *Frontiers in ICT*, 4:29, 2017.
- [NIS97] NIST. Announcing development of a federal information processing standard for advanced encryption standard. <https://csrc.nist.gov/news/1997/announcing-development-of-fips-for-advanced-encryp>, 1997.
- [NIS07] NIST. SHA-3 project. <https://csrc.nist.gov/projects/hash-functions/sha-3-project>, 2007.
- [NIS16] NIST. Submission requirements and evaluation criteria for the post-quantum cryptography standardization process. <https://csrc.nist.gov/CSRC>

/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf, 2016.

- [NIS19] NIST. PQC standardization process: Second round candidate announcement. <https://csrc.nist.gov/News/2019/pqc-standardization-process-2nd-round-candidates>, 2019.
- [OAS12] Frédérique Oggier and B A. Sethuraman. Quotients of orders in cyclic algebras and space-time codes. *Advances in Mathematics of Communications*, 7, 2012.
- [OBV07] Frédérique Oggier, Jean-Claude Belfiore, and Emanuele Viterbo. Cyclic division algebras: A tool for space-time coding. *Foundations and Trends® in Communications and Information Theory*, 4(1):1–95, 2007.
- [OLB09] G Rekaya-Ben Othman, Laura Luzzi, and Jean-Claude Belfiore. Algebraic reduction for the Golden code. In *2009 IEEE International Conference on Communications*, pages 1–5. IEEE, 2009.
- [ORBV06] Frédérique Oggier, Ghaya Rekaya, Jean-Claude Belfiore, and Emanuele Viterbo. Perfect space-time block codes. *IEEE Transactions on information theory*, 52(9):3885–3902, 2006.
- [Pei10] Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In *Annual Cryptology Conference*, pages 80–97. Springer, 2010.
- [Pei16] Chris Peikert. How (not) to instantiate Ring-LWE. In *International Conference on Security and Cryptography for Networks*, pages 411–430. Springer, 2016.
- [Pie82] Richard S. Pierce. Cyclic Division Algebras. In Richard S. Pierce, editor, *Associative Algebras*, Graduate Texts in Mathematics, pages 276–293. Springer New York, New York, NY, 1982.
- [PP19] Chris Peikert and Zachary Pepin. Algebraically structured LWE, revisited. Cryptology ePrint Archive, Report 2019/878, 2019. <https://eprint.iacr.org/2019/878>.

- [PRSD17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of Ring-LWE for any ring and modulus. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 461–473. ACM, 2017.
- [PUTPG⁺19] Alberto Pedrouzo-Ulloa, Juan Ramón Troncoso-Pastoriza, Nicolas Gama, Mariya Georgieva, and Fernando Pérez-González. Revisiting multivariate Ring Learning with Errors and its applications on lattice-based cryptography. Cryptology ePrint Archive, Report 2019/1109, 2019. <https://eprint.iacr.org/2019/1109>.
- [PUTPPG16] Alberto Pedrouzo-Ulloa, Juan Ramón Troncoso-Pastoriza, and Fernando Pérez-González. On Ring Learning with Errors over the tensor product of number fields. *arXiv preprint arXiv:1607.05244*, 2016.
- [Reg05] Oded Regev. On lattices, Learning With Errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):34, 2005.
- [Rei75] I. Reiner. *Maximal orders*. L.M.S. monographs. Academic Press, 1975.
- [Res15] Eric Rescorla. The transport layer security (TLS) protocol version 1.3. <https://tools.ietf.org/html/draft-ietf-tls-tls13-05>, 2015.
- [RNSL17] Martin Roetteler, Michael Naehrig, Krysta M Svore, and Kristin Lauter. Quantum resource estimates for computing elliptic curve discrete logarithms. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 241–270. Springer, 2017.
- [RSA78] Ronald L. Rivest, Adi Shamir, and Len Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.
- [RSW18] Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the Ring-LWE and Polynomial-LWE problems. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 146–173. Springer, 2018.

- [She16] John Sheekey. A new family of linear maximum rank distance codes. *Advances in Mathematics of Communications*, 10(3):475–488, 2016.
- [Sho94] Peter W Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Proceedings 35th annual symposium on foundations of computer science*, pages 124–134. IEEE, 1994.
- [SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 617–635. Springer, 2009.
- [VHLR09] Roope Vehkalahti, Camilla Hollanti, Jyrki Lahtonen, and Kalle Ranto. On the densest MIMO lattices from cyclic division algebras. *IEEE Transactions on Information Theory*, 55(8):3751–3780, 2009.