Université de Rennes 1

Thèse d'Habilitation

On the hardness of the Learning With Errors problem and its variants

présentée et soutenue publiquement le 22 juin 2021 par

Adeline Roux-Langlois
Chargée de Recherche au CNRS

pour obtenir le

Diplôme d'Habilitation à Diriger des Recherches
de l'Université de Rennes 1
Spécialité Informatique

Composition du Jury

Rapporteurs :	Jung Hee Cheon	Professeur, Seoul National University
	Céline Chevalier	Maitre de Conférence HDR, l'Université Paris 2
	Philippe Gaborit	Professeur, l'Université de Limoges
Examinateurs :	Stéphanie Delaune	Directrice de Recherche CNRS, IRISA Rennes
	Maria Naya-Plasencia	Directrice de Recherche INRIA Paris
	David Pointcheval	Directeur de Recherche CNRS, ENS Paris
	Brigitte Vallée	Directrice de Recherche CNRS, GREYC Caen

Contents

Chapter 1

Introduction

The following document is a synthesis of my research activity since my PhD defence in 2014. After one year of post-doc in the LASEC team at EPFL, Switzerland, I joined the EMSEC team in the IRISA laboratory in November 2015. During the past five years, my best achievement was without any doubt my two wonderful children, but also I worked on my research project which was to develop lattice-based cryptography. I was sharing the supervision of 3 PhD students who finished their PhD (Pauline Bert and Chen Qian in 2019, Guillaume Kaim in 2020), and I am currently co-supervising 3 PhD students: Katharina Boudgoust (since 2018), Olivier Bernard (since 2019) and Lucas Prabel (since 2020). After those five years, I am still convinced that lattice-based cryptography is very promising as a viable post-quantum alternative to modern asymmetric cryptography. I am also still enjoying a lot working on this area, which has the particularity to combine theoretical proofs, constructions and real world applications.

1.1 Lattice-based cryptography

Lattice-based cryptography regroups the approaches which consist in building cryptographic constructions and protocols with their security relying on hard problems on lattices. A lattice is defined as a set of all integer linear combinations of some linearly independent vectors that we call a basis. Lattices arise in many different areas of mathematics, such as number theory, geometry and group theory. There are several well-studied hard problems on lattices, such as the Shortest Vector Problem (SVP) which asks to find the shortest non zero vector of a lattice, given a basis, and its approximate variant called Approx-SVP. The security of lattice-based cryptographic constructions relies on an approximate variant of this problem, which depends on an approximation factor γ. This approximate problem is conjectured to be computationally hard to solve when γ is polynomial in the dimension of the lattice. But we need to use intermediate problems to build cryptosystems, such as the Short Integer Solution (SIS) problem [Ajt96], the Learning With Error (LWE) problem introduced by Regev [Reg05], the NTRU problem [HPS98], or one of their many variants. Those problems are fundamentals as they bridge the gap between the security of the cryptographic constructions and the known hard problems on lattices. Indeed, in most constructions, it is possible to build a security proof which ensures that finding an attack against the scheme is at least as hard as solving a hard problem on lattices.

More precisely, for parameters n (the dimension) and q (the modulus), the search version of the Learning With Errors problem asks to find a secret $\mathbf{s} \in \mathbb{Z}_q^n$, given arbitrar-

ily many independent pairs $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e \bmod q)$ where $\mathbf{a}$ is a vector chosen uniformly at random in $\mathbb{Z}_q^n$ and e is an error sampled from a probability density distribution χ. The decision version of LWE consists in distinguishing between arbitrarily many independent pairs $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e \bmod q)$ sampled as just described, and the same number of independently uniformly sampled pairs. In his seminal work, Regev [Reg05] showed that the Learning With Error problem is at least as hard as the decisional variant of approximate SVP by providing a quantum reduction, under certain conditions over the dimension, the modulus and the error distribution χ. He showed that the search and the decision versions of LWE are equivalent, which is particularly interesting in cryptography, as security requirements for encryption schemes, for example, rely on decisional problems which are usually easier than search one. He also described a public key encryption scheme which is proven IND-CPA secure if LWE is hard. This security proof is a reduction, showing that if an attacker succeeds to have a non negligible advantage in the IND-CPA security game, then a challenger can use this advantage to find a solution to the LWE problem. This proof allows to rely the security of the scheme directly on the hardness of the Approx-SVP problem. Nevertheless it is important to notice that in practice, the parameters of a scheme (i.e. related to the dimension, modulus and error distribution) are not chosen to satisfy the conditions of the reduction as it would be too costly in term of efficiency. Schemes parameters are then chosen using the best known attack against the underlying problem, leading to a gap between proven secure constructions and the ones used in practice. Still, the reductions strengthen the confidence we have in the security of the constructions, and provide a valuable help to better understand the underlying difficult problems.

The lattice-based approach to build cryptographic schemes is very promising as we can observe in the ongoing NIST competition for post-quantum cryptography[1]. Indeed, lattice-based cryptography seems to be an interesting candidate to obtain constructions which resist to attacks using a quantum computer. Since the work of Shor [Sho97], the hardness of number theoretical assumptions, which are used as security foundations for many primitives, is extremely reduced when facing a quantum computer. Even if powerful enough quantum computers do not exist yet, it is necessary to be prepared and anticipate their arrival. Over the initial 69 submissions made at the NIST post-quantum competition in 2017, it is worth noticing that 5 over the 7 finalists are related to lattices: the three public key encryption schemes are CRYSTALS-Kyber [BDK+18], NTRU [CDH+18] and SABER [DKRV18, DKR+19], and the two signature schemes are CRYSTALS-Dilithium [DKL+18] and FALCON [FHK+17]. The security of most of those constructions are based on structured variants of LWE [SSTX09, LPR10, LS15] and SIS [LM06, PR06] that we will define later more formally. Those variants, over rings or modules, allow a better efficiency of the schemes by using some algebraic structures.

There are plenty of variants of the LWE problem, and it is important to study them to have a better understanding of their hardness. Indeed, the difficulty of this problem depends on its parameters (the dimension n, modulus q, number of samples m) and on the distribution of its secret (uniform in $\mathbb{Z}_q^n$) and error (following a Gaussian distribution). Indeed, the main drawback of the LWE problem is its efficiency, as we have to use quite large parameters to ensure its hardness. A second possible issue when using LWE is the use of a discrete Gaussian for the error distribution. Indeed, most of the reductions showing the hardness of LWE or its variants need specific properties of Gaussian distributions to work. Recent attacks shows that discrete Gaussian samplers must be built very carefully to avoid side-channel attacks [BHLY16, Pes16, Saa18]. There are several possibilities to avoid the use of Gaussian. Some practical constructions are using other distributions for

[1] `https://csrc.nist.gov/Projects/Post-Quantum-Cryptography`

the error term such as a binomial, by considering that the security is mostly depending on the standard deviation of the distribution. An other idea is to use a deterministic error as in the Learning With Rounding problem (LWR) introduced by [BPR12], which does not depend on Gaussian distributions.

Concerning the efficiency, the idea of using an algebraic structure was first introduced for the SIS problem [LM06, PR06] then used for LWE in [SSTX09] and [LPR10] to define Polynomial LWE and Ring LWE respectively. Those structured variants work on some particular rings instead of over the integers, which allows much more efficient operations (such as the Number Theoretic Transform for instance). The impact on the security is still not clear, as there also exist reductions to show the hardness of those problems, but which seem less powerful than for the original LWE problem. In particular, there is a reduction from Approx-SVP on ideal lattices to Ring LWE [LPR10], where ideal lattices are a class of lattices corresponding to ideals in a specific ring. However, the hardness of Approx-SVP on ideal lattices is still an open problem, and several recent works tend to show that this problem could be easier than the general case (i.e. Approx-SVP on arbitrary lattices), in particular using quantum algorithms. The Module LWE problem was first introduced in [BGV12] and then studied in [LS15]. In this case, a quantum reduction [LS15] showed that Module LWE is at least as hard as Approx-SVP on module lattices. Its interest is to bridge the gap between LWE and Ring LWE (since both can be seen as particular cases of Module LWE) and then to allow a better trade-off between security and efficiency. Variants on Module are indeed used in three finalists of the NIST competition (SABER, CRYSTALS-Kyber and Dilithium).

The NIST competition is a gateway introducing efficient and secure lattice-based schemes. But for now, it only concerns public key encryption, key-exchange mechanism and signature scheme. Applications in cryptography are using a lot additional security properties or more advanced features, and there are still many open problems to build most of those constructions efficiently using lattice assumptions.

1.2 My contributions

In this context, I will present my research work since my PhD in 2014. I will first introduce my contributions on studying the hardness of the Learning With Errors problem and its variants, which is the part I will detail in this manuscript. In particular, Chapter 2 is giving in details the recent state of the art on the hardness of the variants of the LWE problem. I then regroup in Chapter 3 results from 4 research articles.

In a second part of this introduction, I give an overview of an other side of my research work which concerns cryptographic constructions and in particular advanced signature schemes. It concerns 7 research articles, some of them being works in progress.

1.2.1 On studying the hardness of LWE and its variants

In lattice-based cryptography, we consider two types of reductions. The first one are the "worst-case to average-case" reductions, which link the hardness of an "average-case" problem like LWE (as its instances are sampled uniformly at random) to the hardness of a lattice problem for all its instances (including the worst case). The second one are the "average-case to average-case" reductions, which for example are used to show the hardness of LWE with a new error distribution from a LWE with a Gaussian error.

The general ideal behind the notion of quality of a reduction would be that the best algorithm would take the same time to solve both problems. In general, the quality of a

reduction is defined by its tightness. For instance, if we consider a reduction from a problem $\mathcal{P}_2$ to a problem $\mathcal{P}_1$, it means that if we have an algorithm solving $\mathcal{P}_1$ with time T_1 and probability of success ε_1 then the reduction provides an algorithm to solve $\mathcal{P}_2$ with time T_2 and probability of success ε_2 such that $\varepsilon_2 \geq \varepsilon_1$. Informally, a reduction is said to be tight if T_2 is about the same than T_1 and ε_2 is almost equal to ε_1. For example we can consider the ratio $T_2\varepsilon_1/T_1\varepsilon_2$ which should be the closest to 1. In lattice-based cryptography, we also compare the difference of parameters between reductions. Concerning a "worst-case to average-case" reduction from $\mathrm{GapSVP}_{n',\gamma}$ to $\mathrm{LWE}_{n,q,\alpha}$ for example, the quality of the reduction will depend also on the relation between the parameters. If $n' = n$, the reduction will be considered better than if n' is smaller than n. Indeed, if we fix all other parameters the hardness of GapSVP and LWE will increase with their dimension. Concerning "average-case to average-case" reductions, from $\mathrm{LWE}^{m'}_{n',q,\phi_2}$ to $\mathrm{LWE}^{m}_{n,q,\phi_1}$, where ϕ_1 and ϕ_2 are the errors distributions, the best reductions will preserve the number of samples (i.e. $m' = m$) and the dimension (i.e. $n' = n$). Then, we will also compare the parameters of the two distributions ϕ_1 and ϕ_2, in particular the "growth" from ϕ_2 to ϕ_1. If we take all of this in consideration, it explains why it may be hard to compare two reductions for the same result, as one could be samples-preserving compared to the other but with a larger growth between the errors.

I start with the study of the hardness of some variants of the Learning With Errors problem, which is a natural following of my PhD work. I strongly believe that a better understanding of the hardness of those problems, and how they can be linked to one another, is important to reinforce the confidence we need to use them in practice. During my PhD, I first worked on showing the classical hardness of the LWE problem for a polynomial modulus, in a joint work with Zvika Brakerski, Chris Peikert, Oded Regev and Damien Stehlé. I am particularly interested by the Module LWE problem, which I also studied in a joint work with Damien Stehlé [LS15]. Module variants are indeed quite popular, in particular since their use in the NIST competition, as they could be a good compromise to build secure and efficient schemes.

Just after my PhD, in a joint work with Shi Bai, Tancrède Lepoint, Damien Stehlé, Ron Steinfeld and Amin Sakzad, we first introduced the use of the Rényi divergence as a tool to obtain better or easier reductions. We already used the Rényi divergence for a particular application [LSS14] during my PhD with Damien Stehlé and Ron Steinfeld, but we decided to investigate more on how it could be used for other applications. In particular, we showed that it can often be used as an alternative to the statistical distance in security proofs. We proposed several applications such as reducing the storage requirement in the BLISS signature scheme [DDLL13] or obtaining smaller parameters in the Dual-Regev encryption scheme [GPV08]. But in Section 3.1, I focus on two new reductions to study the hardness of variants of LWE. In the first one, we gave an alternative proof to show the hardness of LWE with an uniform error in a small interval, in the second one, we also proposed an alternative proof to study the hardness of the LWR problem.

- [BLL+15, BLR+18] Shi Bai, Tancrède Lepoint, Adeline Roux-Langlois, Damien Stehlé and Ron Steinfeld. *Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance.* Conference version in *ASIACRYPT 2015*, Best Paper Award of the conference, and journal version in the *Journal of Cryptology 2018* with Amin Sakzad as additional author.

Then, with my PhD student Katharina Boudgoust and Weiqiang Wen, and in collaboration with Shi Bai, Dipayan Das and Zhenfei Zhang, we introduced a new variant of

LWR, called Middle Product-LWR, which follows the work of [RSSS17]. The main interest of those middle-product variants is to enjoy a strong security guarantee compared to Polynomial-LWE. In this problem, we work on polynomials in $\mathbb{Z}_q[x]/f$ instead of integers, where $f \in \mathbb{Z}[x]$ is a monic irreducible polynomial. This gain in efficiency comes with a potential decrease in the security. A first issue is that there are concrete examples of polynomials f for which the Polynomial LWE problem becomes easy: for instance when f has a linear factor over $\mathbb{Z}$ [CIV16]. The advantage of middle product problems is that their hardness rely on Polynomial LWE for a set of polynomials instead of only one, meaning that it is sufficient that Polynomial LWE is hard for one of the polynomials in the set for MP-LWE to be hard. Unfortunately, this security is quite costly and for now, the schemes built on middle product problems are not sufficiently efficient.

- [BBD+19] Shi Bai, Katharina Boudgoust, Dipayan Das, Adeline Roux-Langlois, Weiqiang Wen and Zhenfei Zhang. *Middle-Product Learning with Rounding Problem and Its Applications.* In *ASIACRYPT 2019.*

Finally, I present two results on the hardness of Module LWE that we obtained with Katharina Boudgoust, Corentin Jeudy who came doing an internship with us, and Weiqiang Wen. An important open question is to know if one can build a classical reduction to show the hardness of Ring LWE, as it exists for LWE. The only known reduction for Ring LWE for a polynomial modulus is quantum. In our first paper, we study the intermediate case of Module LWE, and show that a classical reduction exists in the linear rank case. It is an interesting first step, even if it does not correspond on instances of the problem used in practical schemes (where the rank is a small constant such as 2 or 3). We hope to be able to generalize it to smaller rank which will be closer to the Ring LWE case but it is still an open problem.

One important step of this reduction is to show the hardness of Module LWE for binary secret (which is also still an open problem for Ring LWE), and in our second result, we study more in details the hardness of this problem. We use a different approach which follows the one in [BLP+13] for LWE, and in particular we study the hardness of the Module Extended-LWE problem. This work allows in particular to improve parameters of the Gaussian noise.

- [BJRW20] Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois and Weiqiang Wen. *Towards Classical Hardness of Module-LWE: The Linear Rank Case.* In *ASIACRYPT 2020.*

- [BJRW21] Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois and Weiqiang Wen. *On the Hardness of Module-LWE with Binary Secret.* In *CT-RSA 2021.*

1.2.2 Building efficient and advanced constructions

I continue with a new part of my research work, on designing and implementing advanced cryptographic constructions. I will not detail this part in the rest of the manuscript, but only explain this line of works in this introduction.

With the recent NIST post-quantum competition, there are several lattice-based candidates for efficient and secure signature and public key encryption schemes, but not for more advanced scheme yet. Indeed, in some cases, such as zero knowledge proofs of knowledges or group signature schemes, there are many nice theoretical propositions, but none

of them achieve the efficiency of constructions currently used. One of the goal of my research regarding this part is to investigate advanced cryptographic constructions, and in particular advanced signature schemes.

To build such advanced constructions, trapdoors are an interesting and very powerful tool. Cryptographic trapdoors allow someone to invert a one way function. It means that without knowing the trapdoor the function is hard to invert, but it becomes easy when you know the trapdoor. First used to build a signature in [GPV08] and developed in several works after [MP12, GM18], lattice trapdoors are very useful, and can be used to build Identity/Attribute Based Encryption [GPV08, ABB10, GVW13, BGG$^+$14], Group Signatures [GKV10, LLLS13]... A lattice trapdoor is defined as a good basis of a particular lattice, allowing to find short elements and especially elements following a Gaussian distribution on the lattice. Unfortunately, the use of trapdoor is still costly in lattice-based cryptography, whether it is to generate the lattice with its trapdoor or to obtain a preimage using this trapdoor. One approach to use trapdoor efficiently is the use of so-called NTRU trapdoors started in [DLP14] and used in the FALCON [FHK$^+$17], which is a NIST signature finalist. In this case, the trapdoor is a good basis of an NTRU lattice, which implies the schemes built using this trapdoor also have to rely on the NTRU assumption.

Implementing lattice-based schemes. I started to work on implementing schemes to test their efficiency just after my PhD. We obtained a first result in 2015, with Martin R. Albrecht, Catalin Cocis and Fabien Laguillaumie, on implementing a candidate multilinear map in 2015. This implementation was the first of a graded encoding scheme over ideal lattices, and one of its main challenges was to handle the size of the parameters induced by such a construction. Multilinear maps (defined as graded encoding schemes) [GGH13a, CLT13, GGH15] were then a popular tool, their main applications was to build a N-party key exchange (which generalizes the 3-party Diffie-Hellman key exchange), or indistinguishability obfuscation [GGH$^+$13b] (iO). But soon after their introduction, a line of cryptanalysis works, starting with the attack of Cheon *et al.* [CHL$^+$15], showed that all these schemes were insecure for most of their applications. When we published this article, the implemented construction from [GGH13a, LSS14] was still secure, and the attacks only concerned the [CLT13] scheme. But soon after, an attack targeting the GGH13 construction was published [HJ16]. After spending two years working with Martin Albrecht on trying to fix those schemes, I decided to stop working on that particular construction.

- [ACLL15] Martin R. Albrecht, Catalin Cocis, Fabien Laguillaumie and Adeline Langlois. *Implementing Candidate Graded Encoding Schemes from Ideal Lattices.* In ASIACRYPT 2015.

One of my first goal was then to obtain a library for implementing lattice-based construction, and in particular trapdoor based ones, that we could use after each time we want to test the efficiency of a new construction. We started this work with an implementation of an identity based encryption (IBE) scheme over rings in 2018, and an associated signature scheme proven secure in the standard model. The goal of this work was also to evaluate if trapdoor based construction could be efficient even without using NTRU lattices. The implementation was made by Pauline Bert and Mohamed Sabt, who were respectively doing a PhD and a post-doc under my supervision. We followed this work by implementing faster trapdoors and adapting them to module in a second work, and then showed that our implementation is modular and can be easily used to implement a

signature scheme and an other IBE. This second implementation was made by Gautier Eberhart (during an internship) and Lucas Prabel. The conclusion of those two works is that those schemes are much slower than the ones using NTRU lattices, but still reasonable if we do not want to rely on the NTRU assumption.

- [BFRS18] Pauline Bert, Pierre-Alain Fouque, Adeline Roux-Langlois and Mohamed Sabt. *Practical Implementation of Ring-SIS/LWE Based Signature and IBE.* In *PQCrypto 2018.*

- Pauline Bert, Gautier Eberhart, Lucas Prabel, Adeline Roux-Langlois and Mohamed Sabt. *Implementation of lattice trapdoors over modules and applications to signatures and IBE.* To appear in *PQCrypto 2021.*

Advanced Signature schemes. The next results come from a collaboration with Sébastien Canard and Jacques Traoré, when supervising the PhD of Guillaume Kaim who was working on lattice-based cryptography for privacy.

Blind signatures, introduced by Chaum [Cha82] in 1982, permit a user to obtain a signature on a chosen message by interacting with a signing authority. We worked on a blind signature scheme, as there was only one existing lattice-based scheme [Rüc10]. Our first goal was to remove the restarts in the Rückert protocol to obtain a practical scheme. But in 2020, Hauck *et al.* [HKLN20] raised a flaw in Rückert security proof, they also proposed a secure construction, with a loss in the efficiency. In our construction, we only managed to partially correct to security flaw in the proof of Rückert, and then, the lattice-based blind signature that we build is not fully proven, and has its security based on a conjecture. Even though we still believe that the conjecture is reasonable, and that this construction could be interesting to use in a practical setting, we decided not to publish this work for now, as the security of the construction is not fully proven yet. Note that conjectures are everywhere in cryptography, and the main question is always how much we believe in them. Finally, thanks to our two aforementioned works on implementing lattice-based schemes, we provided a full implementation of our scheme.

In the following, I present a preliminary result recently accepted to a workshop. Our idea was to use a framework from [FOO92] using blind signatures to build an e-voting scheme. This framework contrasts from the current trend that makes use of homomorphic encryption, or mix-net system [CRS05, CMM19] which are quite efficient and secure thanks to zero-knowledge proofs. However in a post-quantum setting, the lack of efficiency of some of the primitives used in the two frameworks cited above leads us to investigate on new options for a practical e-voting protocol. We still have ongoing works on this construction, such as providing the full security proof, and we also work on an implementation, but we think it could be an interesting alternative to existing post-quantum e-voting schemes which are built using different frameworks [CGGI16, dPLNS17].

- [BCE+20] Samuel Bouaziz-Ermann, Sébastien Canard, Gautier Eberhart, Guillaume Kaim, Adeline Roux-Langlois and Jacques Traoré. *Lattice-based (Partially) Blind Signature without Restart.* IACR Cryptology ePrint Archive 2020: 260 (2020).

- Sébastien Canard, Guillaume Kaim, Adeline Roux-Langlois and Jacques Traoré. *Post-Quantum Resistant E-Voting Scheme.* To appear in *VOTING'21.*

With Guillaume, Sébastien and Jacques, we also worked on group signature schemes, with Adela Georgescu, a post-doc in our team. Group signatures were designed to allow

only members of a group to sign messages on behalf of a group while the identity of the signer remains hidden for the verifier. If necessary, the signature can be opened by an entity called group manager who holds some secret information and reveals the identity of the signer. I already worked on group signatures during my PhD [LLLS13, LLNW14], and today, there is a vast literature concerning group signatures based on lattices. But all of the existing lattice-based group signature schemes are in the random oracle model except the recent construction by Katsumata and Yamada [KY19], which is built without using non interactive zero knowledge proofs. In this work, we used their framework to propose a group signature scheme with forward-security in the standard model. The only pre-existing post-quantum secure forward-secure group signature scheme is built from lattices by Ling et al. [LNWX19] in the random oracle model, following the classical framework of encrypt-then-prove, thus using non-interactive zero-knowledge proofs. Forward-security prevents an attacker in possession of a group signing key to forge signatures produced in the past. In case of exposure of one group member's signing key, group signatures lacking forward-security need to invalidate all group public and secret keys (by re-initializing the whole system) but also invalidate all previously issued group signatures.

- [CGK+20] Sébastien Canard, Adela Georgescu, Guillaume Kaim, Adeline Roux-Langlois and Jacques Traoré. *Constant-size lattice-based group signature with forward security in the standard model.* In ProvSec 2020.

Cryptanalysis. To finish this introduction, I mention a work with Olivier Bernard who is doing a PhD under my supervision and who really wanted to work on cryptanalysis over ideal lattices. Cryptanalysis allows a better understanding of the problems we use, and of the security of the constructions. In lattice-based cryptography it is important to correctly choose the parameters of the cryptographic schemes. Problems on ideal lattices are studied since the result of [SSTX09, LPR10] who showed that there is a reduction from the Approx-SVP on ideal lattices to Ring LWE. For a long time, the best known algorithm to solve this problem on ideal lattices was the same as for arbitrary lattice. But recently, a series of works tends to show that solving this problem could be easier in ideal lattices than in arbitrary ones, in particular using quantum algorithm.

The main contribution of this work is to propose an improved variant of the PHS algorithm [PHS19] with its full implementation which suggests that much better approximation factors are achieved than in the original algorithm.

- [BR20] Olivier Bernard and Adeline Roux-Langlois. *Twisted-PHS: Using the product formula to solve Approx-SVP in ideal lattices.* In ASIACRYPT 2020.

Chapter 2

The Learning with Errors problem

The "Learning With Errors" problem (LWE), first defined by Regev [Reg05], is a fundamental problem in lattice-based cryptography. This problem and its many variants are used to prove the security of a lot of constructions, starting from very basic ones (like public key encryption or signature scheme) to much more advanced (including Fully Homomorphic Encryption).

This chapter starts with the definitions of the Learning With Errors problem in Section 2.1 and a reminder on the basic notions on hard problems on lattices in Section 2.2. Then, I first recall existing results on the hardness of the Learning With Errors problem and its variants. I start with variants associated with other distributions for the secret or the error term in Section 2.3. Section 2.4 then defines the different structured variants, from Polynomial, Ring, Module LWE to the Middle product problem, and gives an overview on the known hardness results concerning those problems. Finally, I recall the deterministic variant of LWE, called the Learning With Rounding (LWR) problem and known results on this problem and its structured variants in Section 2.5.

2.1 Definition of LWE

For q a positive integer, let $\mathbb{T}$ denotes the additive group of reals in $[0,1)$ with addition modulo 1, and $\mathbb{T}_q$ denotes its cyclic subgroup of order q, i.e. $\{0, 1/q, \ldots, (q-1)/q\}$. We denote by ϕ a continuous distribution on $\mathbb{T}$ and by χ a discrete distribution on $\mathbb{Z}$ or $\mathbb{Z}_q$.

The Learning With Errors problem was first defined and studied considering a particular continuous error distribution on $\mathbb{T}$, the main results concerning its difficulty was proven using this distribution. We start by this definition of the problem.

Definition 2.1 ($\mathrm{LWE}_{n,q,\chi}$)**.** *Given parameters n and q, and ϕ a continuous probability distribution on $\mathbb{T}$, and given many independent samples pairs $(\mathbf{a}, \frac{1}{q}\langle \mathbf{a}, \mathbf{s}\rangle + e) \in \mathbb{Z}_q^n \times \mathbb{T}$ where $\mathbf{a}$ is uniformly sampled in $\mathbb{Z}_q^n$ and e is sampled from ϕ, the goal of the search version is to find the secret $\mathbf{s} \in \mathbb{Z}_q^n$.*

An other definition of LWE considers samples of the form $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e) \in \mathbb{T}_q^n \times \mathbb{T}$ where $\mathbf{a}$ is uniformly sampled in $\mathbb{T}_q^n$. Regev showed that it is also possible to use a discrete distribution for the error term, originally by considering a discretization of ϕ, to build cryptographic construction. Then, analogously to the continuous version, it becomes standard to directly define LWE with a discrete error. In this case, we consider a distribution χ on $\mathbb{Z}$ and samples of the form $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s}\rangle + e \bmod q) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$. This last definition seems to be the simplest to use in order to build cryptographic schemes relying on the LWE problem.

The LWE problem also has a decision version, which asks to distinguish between arbitrary many samples of the form $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e \bmod q)$ and uniform samples in $\mathbb{Z}_q^n \times \mathbb{Z}_q$. An interesting aspect of this problem is that the search version of LWE is actually equivalent to its decision version in some cases that we will discuss later [Reg05, Pei09, MP12, BLP+13, MM11]. This equivalence is quite interesting in cryptographic applications, as search variants are usually harder than decision variants. As an example, proving that a public key encryption scheme is IND-CPA secure relies on a decision problem, as we want to show that distinguishing two ciphertexts is hard.

The secret is originally and usually chosen uniformly at random in $\mathbb{Z}_q^n$, but we might also want to use other distributions for the secret as well. This is the case in some reductions, as a smaller secret may be needed, or in some particular applications such as FHE. In this case, we use the notation $\mathrm{LWE}_{n,q,\chi}(\mathcal{D})$ where $\mathcal{D}$ is the distribution of the secret, which is by default uniform in $\mathbb{Z}_q^n$.

For a fixed number of samples m, we use the notation $\mathrm{LWE}_{n,q,\chi}^m$. In this case, we can also define a matrix version of the LWE problem, which asks to find the vector $\mathbf{s}$ given the pair $(\mathbf{A}, \mathbf{As} + \mathbf{e} \bmod q)$, where $\mathbf{A}$ is uniformly sampled in $\mathbb{Z}_q^{m \times n}$, and $\mathbf{e}$ is distributed following the distribution χ^m.

2.2 Preliminaries on lattices

The hardness of the Learning With Errors problem relies on hard problem on lattices such as the decisional version of the Shortest Vector Problem (SVP). We start by defining those problems and recalling why we have a strong confidence on their hardness.

2.2.1 (Hard) problems on lattices

An n-dimensional full-rank lattice Λ is a discrete additive subgroup of $\mathbb{R}^n$. A lattice is the set of all integer combinations of some linearly independent basis vectors, $\mathbf{B} = \{\mathbf{b}_1, \cdots, \mathbf{b}_n\} \in \mathbb{R}^{n \times n}$, $\Lambda(\mathbf{B}) = \{\sum_{i=1}^n z_i \mathbf{b}_i : z_i \in \mathbb{Z}\}$. An important aspect is that a same lattice can be defined by several basis, which has a different quality (see for example Figure 2.1 where the left basis seems better than the one on the right).

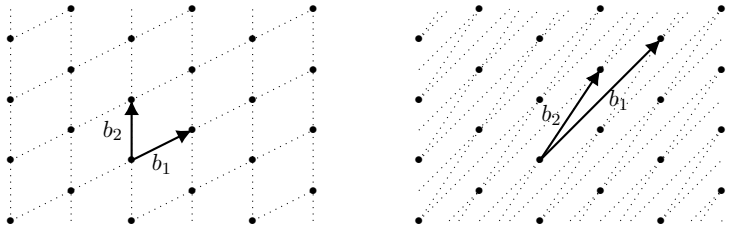

Figure 2.1: Several possible basis for a same lattice.

We define the first minimum of a lattice as $\lambda_1(\Lambda) = \min_{\mathbf{v} \in \Lambda \setminus \{\mathbf{0}\}} \|\mathbf{v}\|_2$ when considering the Euclidean norm. And we define the n-th minima $\lambda_n(\Lambda)$ as the radius of a ball which contains the smallest n linearly independent vectors.

We now define the following well-known problems on lattices, which are the Shortest Vector Problem, its decisional variant GapSVP and the Shortest Independent Vectors problem (SIVP). In cryptography, we use approximate variants of these problems which are defined with an approximation factor $\gamma(n) \geq 1$. We also call the following problems

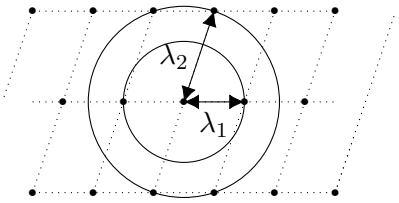

Figure 2.2: Exemple of minima in dimension 2.

Approx-SVP, Approx-GapSVP and Approx-SIVP, when the approximation factor is not specified.

Definition 2.2 (SVP_γ)**.** *Given a n-dimensional lattice Λ, SVP_γ asks to find a non-zero vector of Λ of length less than $\gamma \cdot \lambda_1(\Lambda)$.*

Definition 2.3 (GapSVP_γ)**.** *Given a n-dimensional lattice Λ and $r > 0$, GapSVP_γ asks to decide if $\lambda_1(\Lambda) \leq r$ (YES instance) or $\lambda_1(\Lambda) \geq \gamma \cdot r$ (NO instance) .*

Definition 2.4 (SIVP_γ)**.** *Given a n-dimensional lattice Λ, SIVP_γ asks to find a set of n linearly independent vectors of Λ which all have length less than $\gamma \cdot \lambda_n(\Lambda)$.*

Lattice-based cryptography relies on the conjecture that there is no polynomial algorithm solving those problems, in particular GapSVP_γ, for polynomial approximation factors in the dimension of the lattice. More precisely, SVP is shown to be NP-hard (under randomized reduction) for a constant approximation factors [Ajt98, Mic98, BS99]. Is was also shown in [AR05] that for an approximation factor $\gamma = \sqrt{n}$, GapSVP $\in$ NP$\cap$coNP.

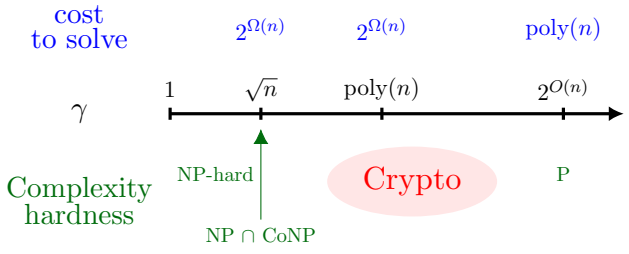

Figure 2.3: Hardness of GapSVP_γ.

The first well known algorithm to solve this problem is the LLL algorithm [LLL82], named after its authors Lenstra, Lenstra and Lovasz. The LLL algorithm allows to find a short vector of an n-dimensional lattice in time $2^{O(n^2)}$ and memory poly(n). Concerning SVP_γ, it also solves this problem for $\gamma = 2^{(n-1)/2}$ in polynomial time. Several algorithms improved this result [Kan83, AKS01, MV10]..., but still resulting in an exponential complexity in the dimension.

2.2.2 Gaussian distributions

Gaussian distributions are used in lattice-based cryptography to study the hardness of fundamental problems but also to build constructions. We recall here some basic definitions on Gaussians. The Gaussian function of center $\mathbf{c} \in \mathbb{R}^n$ and width parameter σ is defined as $\rho_{\sigma,\mathbf{c}}(\mathbf{x}) = \exp(-\pi \frac{\|\mathbf{x}-\mathbf{c}\|^2}{\sigma^2})$, for all $\mathbf{x} \in \mathbb{R}^n$. We also define $D_{\sigma,\mathbf{c}}(\mathbf{x}) = \rho_{\sigma,\mathbf{c}}(\mathbf{x})/\sigma^n$.

Gaussians has two nice properties, first they allow to sample short elements (depending of their parameters) with very high probability, we have in particular for $\mathbf{x} \in \mathbb{R}^n$, that $\Pr_{\mathbf{x} \hookleftarrow D_\sigma}[\|\mathbf{x}\| \geq \sqrt{n}\sigma] \leq 2^{-n}$, and second it is easy to add Gaussian as the sum of two Gaussians of parameters σ and τ results in a Gaussian of parameter $\sqrt{\sigma^2 + \tau^2}$.

The discrete Gaussian distribution over a lattice Λ is obtained by conditioning $D_{\sigma,\mathbf{c}}(\mathbf{x})$ at the event $\mathbf{x} \in \Lambda$, it is defined as $D_{\Lambda,\sigma,\mathbf{c}}(\mathbf{x}) = \frac{D_{\sigma,\mathbf{c}}(\mathbf{x})}{D_{\sigma,\mathbf{c}}(\Lambda)}$, where $D_{\sigma,\mathbf{c}}(\Lambda) = \sum_{\mathbf{x} \in \Lambda} D_{\sigma,\mathbf{c}}(\mathbf{x})$.

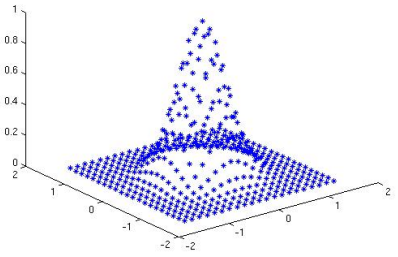

Figure 2.4: Discrete Gaussian distribution on dimension 2

Introduced in 2004 by Micciancio and Regev [MR04], the *smoothing parameter* of a lattice is an important parameter, which intuitively is the minimum parameter for a discrete gaussian distribution on a lattice to behave like a continuous one (in particular concerning the addition and the size properties). It is formally defined as followed.

Definition 2.5. *For all $\varepsilon > 0$, the* smoothing parameter $\eta_\varepsilon(\Lambda)$ *of a lattice Λ with parameter ε is defined as the smallest s such $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$.*

The size of this parameter is linked to the n-th minima of the lattice. In particular, we have [MR04, Reg05] for any n-dimensional lattice Λ and $\varepsilon > 0$,

$$\sqrt{\frac{\ln(1/\varepsilon)}{\pi}} \cdot \frac{\lambda_n(\Lambda)}{n} \leq \eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}} \cdot \lambda_n(\Lambda).$$

It also allows nice properties on discrete Gaussian on lattices such as a bound on the norm of the sampled elements [MR04]. For any n-dimensional lattice Λ, if $s > \eta_\varepsilon(\Lambda)$ then $\Pr_{x \hookleftarrow D_{\Lambda,s,c}}[\|x - c\| > \sqrt{n}s] \leq 2^{-n}$.

2.3 Known results on the hardness of LWE

2.3.1 Quantum and classical reductions

In the following, we now consider that ϕ, the error distribution of LWE, is sampled from a continuous Gaussian of parameter αq, we denote this case by $\mathrm{LWE}_{n,q,\alpha}$. Using a Gaussian distribution is essential in the proofs for the first hardness results of LWE. We discuss further the other possible choices for this error distribution (in particular discrete ones). In the following, the notation LWE stands for the decisional version if nothing is specified.

Quantum reduction. Regev [Reg05, Reg09] was the first to provide a worst-case to average-case reduction showing the hardness of the Learning With Errors problem. This

reduction has the particularity to use a quantum algorithm. This mean given an oracle to solve LWE, we only have a quantum algorithm to solve Approx-GapSVP. Then, if it happens that an attacker can solve LWE, he will need a quantum computer to polynomially solve Approx-GapSVP or SIVP for all possible instances (i.e. even in the worst-case).

Theorem 2.6 ([Reg05, Th. 1.1])**.** *Let $q > 2$ and $\alpha \in (0,1)$ be function of n such that $\alpha q \geq 2\sqrt{n}$. If q is prime and polynomial in n, then there exists a worst-case to average-case quantum polynomial time reduction from* GapSVP_γ *(or* SIVP_γ*) in dimension n to* $\mathrm{LWE}_{n,q,\alpha}$ *for $\gamma = \tilde{O}(n/\alpha)$.*

Regev [Reg05] also showed an important result on the equivalence of the two version of LWE: the decisional and the search ones. It is an interesting property as search problems are usually harder than decision ones, but in security proofs we often rely on a decisional problem. The first direction of the reduction is quite direct since if we are able to solve the search problem, the resulting secret would allow to distinguish between an LWE distribution and a uniform one.

Regev also showed the other direction, that if one can solve the decisional version he can also solve the search version. The idea of this reduction is to find one by one each coordinate of the secret using the oracle for decisional LWE. This first reduction is provided only for q prime and $\mathrm{poly}(n)$, but for any error distribution. It also needs $m' = \tilde{O}(nmq/\varepsilon^2)$ search LWE samples to call the decision oracle with m samples, where ε is the advantage of the oracle to distinguish between the two distributions. The conditions on the parameters of this reduction (in particular on q) was later improved in several works [Pei09, MP12, BLP+13], but only for Gaussian distributions, and with a similar loss on the number of samples. There also exists a sample-preserving search-to-decision reduction given by Micciancio and Mol [MM11], this reduction only works for particular cases, as for instance q prime and $\mathrm{poly}(n)$ for any discrete error on $\mathbb{Z}_q$.

Finally, Regev also provided a worst-case to average-case reduction for LWE using a simple transformation (i.e. adding $\langle \mathbf{a}, \mathbf{t} \rangle$ to $\mathbf{b}$ for an uniform $\mathbf{t}$ to transform a LWE sample with a secret $\mathbf{s}$ to a LWE sample with a secret $\mathbf{s} + \mathbf{t}$ and an uniform sample to an uniform one). This reduction means that if one has a distinguisher for a non negligible fraction of possible secret $\mathbf{s}$ then it is possible to build a distinguisher for all secret $\mathbf{s}$.

First classical reduction. In 2009, Peikert [Pei09] gave the first classical reduction to show the hardness of the LWE problem, but his result was limited to an exponential modulus q in the dimension n. Compared to the original reduction of Regev, this reduction is also limited to GapSVP.

Theorem 2.7 ([Pei09, Th. 3.1])**.** *Let $q > 2$ and $\alpha \in (0,1)$ be function of n, there exists a worst-case to average-case classical polynomial time reduction from* GapSVP_γ *in dimension n to* $\mathrm{LWE}_{n,q,\alpha}$ *for $\gamma \geq \frac{n}{\alpha \log n}$ and $q \geq 2^{n/2} \cdot \omega(\sqrt{\log n / n})$.*

Classical hardness for a polynomial modulus. In 2013 [BLP+13], during my PhD and in collaboration with Zvika Brakerski, Chris Peikert, Oded Regev and my supervisor Damien Stehlé, we provided a new classical reduction to show the hardness of the Learning With Errors problem, for a polynomial modulus q in the dimension n. This smaller size of the modulus is the one used in practice, our reduction also remove the condition on q to be prime.

Theorem 2.8 ([BLP+13])**.** *Let $q > 2$ and $\alpha \in (0,1)$ be function of n such that $\alpha q \geq \sqrt{n}$. There exists a worst-case to average-case classical polynomial time reduction from* GapSVP_γ *in dimension $\sqrt{n}$ to* $\mathrm{LWE}_{n,q,\alpha}$ *for $\gamma = \tilde{O}(n^2/\alpha)$.*

The main idea of this reduction is to start from Peikert's result, i.e. a classical reduction from GapSVP to LWE for an exponential modulus. And then to use a modulus switching technique to provide a reduction from LWE with an exponential modulus to LWE with a polynomial modulus, at the expense of making the error term grows. Note than an important step for this reduction to work is to use LWE with a smaller secret. Overall, this reduction looses a factor $\sqrt{n}$ in the dimension of the two problems, if we compare to Regev's quantum reduction. Those results can be summarised in Table 2.1.

	dim n'	q	α	γ	
[Reg05]	n	poly(n), prime	$\alpha q \geq 2\sqrt{n}$	$\tilde{O}(n/\alpha)$	quantum
[Pei09]	n	$q > 2^{n/2}$	$\alpha q \geq n$	$\geq n/(\alpha \log n)$	classical
[BLP$^+$13]	$\sqrt{n}$	poly(n)	$\alpha q \geq \sqrt{n}$	$\tilde{O}(n^2/\alpha)$	classical

Table 2.1: Reductions from $\mathrm{GapSVP}_{n',\gamma}$ to $\mathrm{LWE}_{n,q,\alpha}$.

Finally, an important consequence of the work of [BLP$^+$13] is to show that the hardness of LWE in dimension n and with a modulus q is a function of $n \log q$. This is shown by the modulus switching reduction which gives a reduction from $\mathrm{LWE}_{n,q}$ to $\mathrm{LWE}_{n/k,q^k}$ and from LWE_{n,q^k} to $\mathrm{LWE}_{nk,q}$.

2.3.2 Possible distributions of the secret

In the original LWE problem, the secret is uniformly sampled in $\mathbb{Z}_q^n$. But there are other possible distributions for the secret, which were introduced in several articles for different cryptographic applications.

Gaussian secret. In 2009, the first result [ACPS09] by Applebaum, Cash, Peikert and Sahai showed that the secret could have the same distribution as the error. In particular, for q a prime power, they provided a polynomial-time reduction from $\mathrm{LWE}_{n,q,\chi}(U(\mathbb{Z}_q^n))$ to $\mathrm{LWE}_{n,q,\chi}(\chi^n)$. As a consequence, the LWE problem with both secret and error following a discrete Gaussian distribution can also be considered as a hard problem, and is usually called HNF-LWE. This reduction was extended to any q in [BLP$^+$13].

Binary secret. In 2010, Goldwasser, Kalai, Peikert and Vaikuntanathan [GKPV10] studied the robustness of the Learning With Errors assumption by considering leaky secrets. They considered that the secret follows a distribution $\mathcal{D}$ with min-entropy k, and showed that there exists a polynomial time reduction from $\mathrm{LWE}_{\ell,q,\alpha}(U(\mathbb{Z}_q^n))$ to $\mathrm{LWE}_{n,q,\beta}(\mathcal{D})$ for $k \geq \ell \log q + \omega(\log n)$ and $\alpha/\beta = \mathrm{negl}(n)$. The main drawback of their result is this last condition on the parameters, which comes from a noise flooding argument in the reduction, as β needs to be very large to satisfy it. This gives in particular the first reduction to show the hardness of the binary variant of LWE, called bin-LWE, where the secret is sampled uniformly in $\{0,1\}^n$.

We improved this result in [BLP$^+$13], as we needed for our main reduction to use the hardness of LWE with binary secret, but with a better bound between the Gaussian parameters. The high level idea of our reduction is similar to the one in [GKPV10], but to avoid the last argument using noise flooding, we used and proved the hardness of the extended-LWE problem. The principle of this new problem is to give an additional information on the noise of the LWE samples. This allows to achieve a reduction from $\mathrm{LWE}_{\ell,q,\alpha}$

to bin-LWE$_{n,q,\beta}$ for $n \geq \ell \log q + \omega(\log n)$ and $\alpha/\beta = 1/\sqrt{10n}$. In 2018, Micciancio [Mic18] proposed a much simpler reduction to prove the hardness of LWE with binary secret, with a similar conditions on the parameters, i.e. $\alpha/\beta = 1/(2\sqrt{n+1})$.

Entropic secret. Recently, Brakerski and Döttling [BD20] extended the hardness of LWE to more general secret distribution. They showed that the hardness depends directly on a property of the distribution of the secret that they call the noise lossiness. One difference with the previous results is that they loose a factor $\sqrt{m}$, where m is the number of samples, between the Gaussian parameters.

2.3.3 Possible distributions of the error

The original quantum reduction of Regev [Reg05] and the two classical reductions [Pei09, BLP+13] all use a continuous gaussian distribution for the error term of the LWE sample. The distribution defined by Regev is a distribution on $\mathbb{T}$ called ψ_β for a parameter $\beta \in \mathbb{R}^+$, which is obtained by sampling from a normal variable with standard deviation $\beta/\sqrt{2\pi}$ and reducing the result modulo 1. Regev also showed in [Reg05, Lemma 4.3] that the problem is still hard using a discretization of the probability density function for the error on $\mathbb{Z}_q$ (and in particular in the case of rounded Gaussian).

Discrete Gaussian. Discrete Gaussian directly sampled from $D_{\mathbb{Z},\alpha q}$ for $\alpha \in [0,1)$ is a popular distribution for the error for many of the applications. The reduction from LWE with a continuous Gaussian distribution to LWE with a discrete one is formalized by Peikert's result [Pei10] using Theorem 3.1. This article also proposes an efficient sampler for discrete Gaussian distributions over a lattice. However building efficient and secure Gaussian samplers appear to be a complicated task, and recent results show that it is possible to use side-channel attacks against existing samplers [BHLY16, Pes16, Saa18]. An other drawback of Gaussian distributions is that the existing sampler are not exacts ones, then provide different distributions than the exact one used in the reduction. Then, and mainly for practical reasons, other distributions were studied for the error term of LWE.

Uniform noise. The LWE problem with noise uniform in a small interval $[-\beta, \beta]$ was first introduced by Döttling and Müller-Quade [DM13]. The authors exhibit a reduction from LWE with Gaussian noise. The main proof ingredients are the construction of lossy codes for LWE (which are lossy for the uniform distribution in a small interval), and the fact that lossy codes are pseudorandom. The reduction from [DM13] needs the number of LWE samples to be fixed and bounded by poly(n), and $\beta \geq mn^c\alpha$ where α is the LWE Gaussian noise parameter and $c \in (0,1)$ is an arbitrarily small constant. Note that it also degrades the LWE dimension by a constant factor, but that there is no condition on q.

The LWE problem with uniform noise in a small interval is also investigated in [MP13], with a focus on the number of LWE samples. The reduction from [MP13] does not preserve the LWE dimension either. Another reduction for LWE with uniform noise can be obtained by using the hardness result for the Learning With Rounding (LWR) problem (defined in Section 2.5) from [BGM+16]. The resulting reduction maps the LWE$_{n',q,\alpha,m}$ problem to the LWE$_{n,q,U([-\beta,\beta]),m}$ problem with $n' = n/\log q$ and $\beta = \Omega(m\alpha/\sqrt{\log n})$. We compare those results with one of our results in Section 3.1.

Binary noise. The particular case of a binary noise is introduced in [MP13]. Micciancio and Peikert showed the hardness of this problem, by providing a reduction from LWE,

but with a bounded number of samples $m = n \cdot (1 + \Omega(1/\log n))$. Indeed, the problem becomes insecure for more samples, in particular using Arora-Ge attack as studied in [AG11, BGPW16].

2.4 Structured variants of LWE

The first drawback of using directly the LWE problem is the lack of efficiency of the constructions that we can build. Indeed, the matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is typically part of the public key and has size $nm \log q$ bits with quite large n and m. To improve this efficiency, many structured variants of LWE was introduced starting with Polynomial LWE (PLWE) [SSTX09] and Ring LWE (RLWE) [LPR10]. The problems described in this section are the one we want to use in practice, as they are the only ones allowing to achieve a sufficient efficiency for most of the cryptographic constructions. But there is no consensus yet on which one would be the better choice.

2.4.1 Polynomial, Ring and Module LWE

As we will keep a high level defining the problems and the results, we omit the algebraic number theories preliminaries for which we refer the reader to [BJRW20, Section 2.1]. We start by defining Polynomial, Ring and Module LWE:

- **Polynomial LWE.** In search PLWE, instead of considering elements of $\mathbb{Z}^n$, we consider elements of $\mathbb{Z}[x]/f$ for a monic irreducible polynomial $f \in \mathbb{Z}[x]$. Given arbitrarily many samples $(a_i, a_i \cdot s + e_i)$ where a_i is uniformly sampled in $\mathbb{Z}_q[x]/f$ and e is small, sampled from a distribution in $\mathbb{Z}[x]/f$, the goal is to find the secret $s \in \mathbb{Z}_q[x]/f$.

- **Ring LWE.** In search RLWE, we consider K a number field of degree n and $\mathcal{O}_K$ its ring of integers. Given arbitrarily many samples $(a_i, a_i \cdot s + e_i)$ where a_i is uniformly sampled in $\mathcal{O}_K/q\mathcal{O}_K$ and e is small, sampled from a distribution in $K_\mathbb{R} := K \otimes_\mathbb{Q} \mathbb{R}$, the goal is to find the secret s. Depending on the variant (primal [DD12] or dual [LPR10]), the secret can be either in $\mathcal{O}_K/q\mathcal{O}_K$ or in its dual $\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee$ (A reduction from dual-RLWE to primal-RLWE is given in [DD12]).

- **Module LWE.** In search MLWE, first mentioned in [BGV12], the goal of this problem is to bridge the gap between LWE and RLWE to obtain a better trade-off between security and efficiency. Given $(\mathbf{a}, b = \langle \mathbf{a}, \mathbf{s} \rangle + e)$, where $\mathbf{a}$ is uniformly sampled in $U((\mathcal{O}_K/q\mathcal{O}_K)^d)$ and e is small, sampled from a distribution in $K_\mathbb{R}$, the goal is to find the secret vector $\mathbf{s} \in (\mathcal{O}_K^\vee/q\mathcal{O}_K^\vee)^d$. We call the parameter d the rank of the module, and the special case of $d = 1$ is exactly RLWE.

The first result on the hardness of the Polynomial LWE problem was given by Stehlé, Steinfeld, Tanaka and Xagawa [SSTX09], which provided a quantum polynomial-time reduction from SVP_γ but restricted on ideal lattices, to search PLWE. Note that this restriction to ideal lattices, which is a class of lattices corresponding to ideals in $\mathbb{Z}_q[x]/f$, is important here and will be discussed later.

In [LPR10], Lyubashevsky, Peikert and Regev then studied the dual variant and provided a quantum polynomial-time reduction from SVP_γ, for $\gamma = \tilde{O}(\sqrt{n}/\alpha)$, restricted to the class of Euclidean lattices corresponding to ideals of $\mathcal{O}_K$, to search RLWE if $\alpha q \geq \omega(\sqrt{\log n})$. They also provided an important search-to-decision reduction when K is a cyclotomic and the modulus q is polynomial in n, prime and such that $q = 1 \bmod 2n$.

The nice case of power-of-two cyclotomic is a very popular choice, in particular using the ring $R = \mathbb{Z}[x]/\langle x^n + 1\rangle$ for n a power of 2. In this case, primal-RLWE and PLWE are the same problem, and as $R^\vee = \frac{1}{n}R$, the dual and primal variants of RLWE are equivalents.

In 2015, during my PhD and with Damien Stehlé [LS15], we studied the Module Learning With Errors problem (MLWE) and provided a quantum reduction from Module SIVP_γ (SIVP restricted to module lattices) to MLWE. We also showed that this reduction admits a converse one, from MLWE to Module SIVP. This is also the case for the LWE problem but not for Ring LWE. We also showed that Ring LWE is hard for all sufficiently large q, independently of the arithmetic properties of q. This result is obtained by adapting the modulus switching technique described in [BLP+13].

In 2017, Albrecht and Deo [AD17] provided a reduction from MLWE in rank d and modulus q to MLWE in rank d/k and modulus q^k in the case where R is a power-of-two cyclotomic ring, and the same for both problems. A consequence is that if we consider $k = d$, there is a reduction from search MLWE in rank d and modulus q to search RLWE with a modulus q^k. This last reduction has an error rate expansion for the Gaussian distribution from α to $\alpha \cdot n^2\sqrt{d}$. The consequence of this work is then that the Ring LWE problem in dimension n and an exponential modulus q^d is hard under the hardness of Approx-SIVPγ on modules.

In 2017, Peikert, Regev and Stephens-Davidowitz [PRS17] gave a direct quantum reduction from ideal lattice problems directly to the decision version of dual-RLWE. The main advantage of their reduction is that it works for any number field and any modulus. The general case is then studied by Rosca, Stehlé and Wallet [RSW18], who showed that the decision and search versions of dual and primal RLWE and PLWE reduce to one another in polynomial time, with a limited error increase and for "huge classes of rings".
Remark. The reductions to the decisional LWE problem on rings or modules is using a particular distribution for the error which is not exactly a Gaussian and is called Υ, it is formally defined in [LPR10, PRS17].

2.4.2 Known weaknesses of structured variants

Gaining in efficiency on the positive side comes with a potential decrease in the security level guarantees on the negative side. There are in particular two lines of work discussing the weaknesses of using Ring/Polynomial LWE. First, there are concrete examples of polynomials f for which the PLWE becomes computationally easy: for instance when f has a linear factor over $\mathbb{Z}$ [CIV16]. Note that this case is excluded by restricting to irreducible polynomials. A review on the known weak instances of PLWE and RLWE is given by Peikert [Pei16]. In his work, Peikert clearly explains that those attacks do not work on instances which satisfy the hypothesis of the "worst-case to average-case" theorems.

The second question concerns the hardness of solving Approx-SVP on ideal lattices. Approx-SVP is a well-known hard problem on lattices (see Definition 2.2), which asks to find short vectors on a given lattice, but its variant restricted to ideal lattices (corresponding to ideals of the ring of integers R of a number field K) is still not fully understood. For a long time, the best known algorithm to solve this problem on ideal lattices was the same as for arbitrary lattices. The best trade-off in this case is given by Schnorr's hierarchy [Sch87], which allows to reach an approximation factor $2^{\tilde{O}(n^\alpha)}$ in time $2^{\tilde{O}(n^{1-\alpha})}$, for $\alpha \in (0, 1)$, using the BKZ algorithm [SE94]. As already mentioned in the introduction, a recent line of works [CGS14, EHKS14, BS16, CDPR16, CDW17, DPW19, PHS19] tends to show that solving this problem could be easier in ideal lattices than in arbitrary ones, in

particular in the quantum setting. Concerning those works, it seems important to remember that there is no converse reduction from Ring LWE to Approx-SVP over ideals. Hence a quantum (or classical) algorithm solving Approx-SVP on ideal lattices would only mean that the reduction from Approx-SVP on ideals to Ring LWE is not an evidence of the hardness of Ring LWE any more, but it would not mean that Ring LWE is easy to solve.

2.4.3 Middle Product LWE

Motivated by the question of how to choose a good polynomial, Lyubashevsky introduces the so-called $R^{<n}$-SIS problem [Lyu16], a variant of the *Short Integer Solution* (SIS) problem, whose hardness does not depend only on a *single* polynomial, but on a set of polynomials. Inspired by this, Roşca, Sakzad, Stehlé and Steinfeld [RSSS17] propose its LWE counterpart: the *Middle-Product Learning With Errors* (MP-LWE) problem.

The MP-LWE problem is defined as follows: Taking two polynomials a and s over $\mathbb{Z}_q$ of degrees less than n and $n+d-1$, respectively, the middle-product $a \odot_d s$ is the polynomial of degree less than d given by the middle d coefficients of $a \cdot s$. In other words, $a \odot_d s = \lfloor (a \cdot s \bmod x^{n+d-1})/x^{n-1} \rfloor$, where the floor rounding $\lfloor \cdot \rfloor$ denotes deleting all those terms with negative exponents on x. Instead of choosing a and s from the ring $\mathbb{Z}_q[x]/f$ as in the PLWE setting, they are now elements of $\mathbb{Z}_q^{<n}[x]$ and $\mathbb{Z}_q^{<n+d-1}[x]$. Here, $\mathbb{Z}_q^{<n}[x]$ denotes the set of all polynomials with coefficients in $\mathbb{Z}_q$ of degree less than n for $n \geq 1$. For integers d, n and q with $q \geq 2$ and $0 < d \leq n$ as parameters, an MP-LWE sample is given by $(a, b = a \odot_d s + e \bmod q)$, where s is a fixed element of $\mathbb{Z}_q^{<n+d-1}[x]$, a is sampled from the uniform distribution over $\mathbb{Z}_q^{<n}[x]$ and e is sampled from a probability distribution χ over $\mathbb{R}^{<d}[x]$.

Regarding the hardness of MP-LWE, Roşca et al. [RSSS17] established a reduction from the PLWE problem parametrized by a polynomial f to the MP-LWE problem defined independently of any such f. Thus, as long as the PLWE problem defined over some f (belonging to a huge family of polynomials) is hard, the MP-LWE problem is also guaranteed to be hard.

2.5 Learning With Rounding

As we already discussed, a second drawback of using the LWE problem in practice is when we want to sample the error, in particular when it is following a discrete Gaussian distribution. It is a natural choice in theory but does not seem to be in practice.

To avoid this problem, an other idea is to consider the Learning With Rounding problem (LWR). Introduced by Banerjee, Peikert and Rosen [BPR12] in 2012, this variant of LWE uses a deterministic error, which corresponds to a rounding of the second term of the LWE pair. For $q \geq p \geq 2$, an LWR sample has the form $(\mathbf{a}, \lfloor \langle \mathbf{a}, \mathbf{s} \rangle \rceil_p) \in \mathbb{Z}_q^n \times \mathbb{Z}_p$, where $\mathbf{a}$ is chosen uniformly at random and $\lfloor \cdot \rceil_p : \mathbb{Z}_q \to \mathbb{Z}_p$ defines a rounding function such that for $x \in \mathbb{Z}_q$, $\lfloor x \rceil_p = \lfloor \frac{p}{q} x \rceil \bmod p$. The LWR problem is then asking, given arbitrary many samples, to find the secret $\mathbf{s}$.

Similarly to LWE, the ring variant of LWR is considering elements on the ring R of dimension n instead of $\mathbb{Z}^n$, and its module variant considers elements on R^d for a rank d.

Hardness of LWR. The following state of the art on the hardness of LWR and its ring version comes from [BBD+19]. In the full version of their paper, published on the IACR Cryptology ePrint Archive, Banerjee et al. [BPR11] showed a reduction from LWE

to LWR with arbitrarily many samples, which also works for the ring counterpart. Unfortunately, the reduction requires q/p to be larger than the error size B (where B bounds the magnitude of the LWE error with high probability) by a super-polynomial factor, thus leading to a large modulus paired with a small error bound. This in turn implies that the underlying worst-case lattice-problems are assumed to be hard with super-polynomial approximation factors, which stands for a stronger assumption. In practice, this also leads to inefficient protocols.

Subsequent studies proposed new reductions that work for a larger range of parameters. Alwen et al. [AKPW13] gave a reduction that allows a polynomial modulus and modulus-to-error ratio. However, the modulus q in the reduction depends on the number of LWR samples, thus the number of samples needs to be fixed in prior by some polynomials. Furthermore, the reduction imposes certain number theoretical restrictions on the modulus q. For example, power-of-two moduli are not covered. In a recent work, Bogdanov et al. [BGM$^+$16] used the Rényi divergence to show a sample preserving reduction, which is also dimension preserving for the special case of prime modulus. They also provided a reduction from the search variant of RLWE to the search variant of RLWR. In another work, Alperin-Sheriff and Apon [AA16] further improved the parameter sets for the reduction. In particular, their reduction is dimension-preserving with a polynomial-sized modulus.

Ring setting. The original work of [BPR12] also considered the ring setting, but with the drawback that the ratio q/p needs to be super-polynomial in the decisional variant. In the search variant, the Ring LWR problem is proven hard thanks to the use of the Rényi divergence. Nevertheless, due to the simplicity and efficiency of RLWR, several schemes as `SABER` [DKR$^+$19] and `Round5` [BBF$^+$19] basing their hardness on RLWR or MLWR are currently participating in the NIST standardization process.

To overcome the lack of provable hardness for decisional RLWR with practical parameters, in 2018, Chen et al. [CZZ18] proposed a new assumption, called the *Computational Learning With Rounding Over Rings* (R-CLWR) assumption. They showed a reduction from decisional RLWE to R-CLWR, where the secret in the RLWE sample is drawn uniformly at random from the set of all invertible ring elements whose coefficients are small.

Finally, a reduction from search to decision RLWR was recently given in [LW20], but considering a new notion for the rounding operation in the ring.

Chapter 3

Recent contributions on the hardness of LWE and its variants

As mentioned in the introduction, the first part of my research work is on a better understanding of the hardness of the Learning With Errors problem and its variants. I am interested, in particular, by reductions which are very powerful to link problems hardness to one another.

In this chapter, I first recall the notion of Rényi divergence in Section 3.1, and explain how it can be used to improve reductions and hence security proofs in cryptography (with an example in Section 3.1.2). Then, I describe in Section 3.2.1 two reductions showing the hardness of the Module Learning With Error problem with a binary secret, the first one using the Rényi divergence and the second one adapting the proof in [BLP+13]. In Section 3.2.2, I explain our classical "worst-case to average-case" reduction to show the hardness of Module LWE when the rank is linear in the dimension. Finally, in Section 3.3, I describe two results on the hardness of LWR, the first one directly using the Rényi divergence on a reduction from LWE to LWR, and the second one by introducing a new problem, called Middle Product Computational Learning With Rounding.

The goal of this chapter is to provide a high level overview of my results, the full version of each proof is available in the corresponding published article [BLL+15, BLR+18, BBD+19, BJRW20, BJRW21], all publicly available on the IACR eprint archive.

3.1 On the use of Renyi divergence

We started this work on the applications of Rényi divergence just after my PhD with Shi Bai, Tancrède Lepoint, Damien Stehlé and Ron Steinfeld, later joined by Amin Sakzad. Our starting point was the use of the Rényi divergence in [LSS14], and we wanted to study the impact of using this Rényi divergence instead of statistical distance in different applications. There were several results on security proofs in this work where we showed several independent applications, I choose to highlight two of them related to reductions. In the following part, I start with a reduction from LWE with a Gaussian noise to LWE with a small uniform noise which is published in [BLL+15].

This work has been awarded at the Asiacrypt 2015 conference and the Rényi divergence has been used in many publications since 2015 (google scholar is recording around 130 citations in March 2021). Rényi divergence is used to propose new reductions or improve existing ones, a good example is that two of the three papers that I present in this chapter makes use of it. It is also used for more practical applications, as shown in the original

paper with two examples. In particular, it is used in some of the NIST submissions to study the distance between a sampled distribution and the theoretical one.

3.1.1 Rényi divergence and security proofs

Let D_1 and D_2 be two non-vanishing probability distributions over a common measurable support X. Let $a \in (1, +\infty)$, the *Rényi divergence* [Rén61, vEH14] $R_a(D_1\|D_2)$ of order a between D_1 and D_2 is defined as $R_a(D_1\|D_2) = \left(\sum_{x\in \mathrm{Supp}(D_1)} \frac{D_1(x)^a}{D_2(x)^{a-1}}\right)^{\frac{1}{a-1}}$, it is the ($(a-1)^{\text{th}}$ root of the) expected value of $(D_1(x)/D_2(x))^{a-1}$ over the randomness of x sampled from D_1. For notational convenience, our definition of the Rényi divergence is the exponential of the classical definition [vEH14]. The Rényi divergence is an alternative to the statistical distance $\Delta(D_1, D_2) = \frac{1}{2}\sum_{x\in X} |D_1(x) - D_2(x)|$ as a measure of distribution closeness, where we replace the difference in the statistical distance, by the ratio in Rényi divergence. The Rényi divergence enjoys several properties that are analogous of those enjoyed by the statistical distance, except that the additive property of the statistical distance is replaced by a multiplicative property in Rényi divergence.

The statistical distance is ubiquitous in cryptographic security proofs. One of its most useful properties is the so-called *probability preservation property*: For any measurable event $E \subseteq X$, we have $D_2(E) \geq D_1(E) - \Delta(D_1, D_2)$. The Rényi divergence enjoys the analogous property $D_2(E) \geq D_1(E)^{\frac{a}{a-1}}/R_a(D_1\|D_2)$. If the event E occurs with significant probability under D_1, and if the statistical distance (resp. Renyi divergence) is small, then the event E also occurs with significant probability under D_2. These properties are particularly handy when the success of an attacker against a given scheme can be described as an event whose probability should be non-negligible, e.g., the attacker outputs a new valid message-signature pair for a signature scheme. If the attacker succeeds with good probability in the real scheme based on distribution D_1, then it also succeeds with good probability in the simulated scheme (of the security proof) based on distribution D_2.

To make the statistical distance probability preservation property useful, it must be ensured that the statistical distance $\Delta(D_1, D_2)$ is smaller than any $D_1(E)$ that the security proof must handle. Typically, the quantity $D_1(E)$ is assumed to be greater than some success probability lower bound ε, which is of the order of $1/\mathrm{poly}(\lambda)$ where λ refers to the security parameter. As a result, we must have the statistical distance $\Delta(D_1, D_2) < \varepsilon$ for the statistical distance probability preservation property to be relevant. In contrast, the Rényi divergence probability preservation property is non-vacuous when the Rényi divergence is $R_a(D_1\|D_2) \leq \mathrm{poly}(1/\varepsilon)$. In many cases, the latter seems less demanding than the former: in all our applications, the Rényi divergence between D_1 and D_2 is small enough for the Rényi divergence probability preservation property while their statistical distance is too large for the statistical distance probability preservation to be applicable.

Although the Rényi divergence seems more amenable than the statistical distance for search problems, it seems less so for distinguishing problems. A typical cryptographic example is semantic security of an encryption scheme. Semantic security requires an adversary $\mathcal{A}$ to distinguish between the encryption distributions of two plaintext messages of its choosing: the distinguishing advantage $\mathrm{Adv}_{\mathcal{A}}(D_1, D_2)$, defined as the difference of probabilities that $\mathcal{A}$ outputs 1 using D_1 or D_2, should be sufficiently large. In security proofs, algorithm $\mathcal{A}$ is often called on distributions D_1' and D_2' that are close to D_1 and D_2 (respectively). If the statistical distances between D_1 and D_1' and D_2 and D_2' are both bounded from above by ε, then, by the statistical distance probability preservation property (used twice), we have $\mathrm{Adv}_{\mathcal{A}}(D_1', D_2') \geq \mathrm{Adv}_{\mathcal{A}}(D_1, D_2) - 2\varepsilon$. As a result, the statistical

distance can be used for distinguishing problems in a similar fashion as for search problems. The multiplicativity of the Rényi divergence probability preservation property seems to prevent Rényi divergence from being directly applicable to distinguishing problems.

3.1.2 Hardness of LWE with small uniform noise.

One of our application is to provide an alternative proof that the Learning With Errors problem with noise chosen uniformly in an interval is no easier than the Learning With Errors problem with Gaussian noise [DM13]. Our reduction is marginally more general as it also applies to distributions with smaller noises. Moreover, our reduction preserves the dimension n of LWE, and is hence tighter than the one from [DM13] (which degrades the LWE dimension by a constant factor).

Theorem 3.1. *Let $\alpha, \beta > 0$ be real numbers with $\beta = \Omega(m\alpha/\log n)$ for positive integers m and n. Let $m > \frac{n \log q}{\log(\alpha+\beta)^{-1}} \geq 1$ with $q \leq poly(m, n)$ prime. Then there is a polynomial-time reduction from $\mathrm{LWE}_{n,q,D_\alpha,m}$ to $\mathrm{LWE}_{n,q,\phi,m}$, with $\phi = \frac{1}{q}\lfloor qU_\beta \rceil$.*

The reduction relies on five steps that are fully detailled in [BLR+18, Section 5].

- A reduction from $\mathrm{LWE}_{n,q,D_\alpha,m}$ to $\mathrm{LWE}_{n,q,\psi,m}$ with $\psi = D_\alpha + U_\beta$, which is quite direct by adding noise to samples.
- A reduction from $\mathrm{LWE}_{n,q,\psi,m}$ to search $\mathrm{LWE}_{n,q,\psi,m}$, which is also quite direct as a decision-to-search reduction.
- A reduction from search $\mathrm{LWE}_{n,q,\psi,m}$ to search $\mathrm{LWE}_{n,q,U_\beta,m}$, using that the Rényi divergence $R_2(U_\beta||\psi)$ can be bounded by $1 + 1.05 \cdot \frac{\alpha}{\beta}$.
- A reduction from search $\mathrm{LWE}_{n,q,U_\beta,m}$ to search $\mathrm{LWE}_{n,q,\phi,m}$, with $\phi = \frac{1}{q}\lfloor qU_\beta \rceil$, which is a reduction from continuous to discrete noise.
- A reduction from search $\mathrm{LWE}_{n,q,\phi,m}$ to $\mathrm{LWE}_{n,q,\phi,m}$, using the [MM11] search-to-decision reduction (as q is prime).

We remark that the search-decision equivalence idea in the proof of Theorem 3.1 could be extended to show the hardness of the decision LWE problem with any noise distribution ψ, with respect to the hardness of LWE with Gaussian noise D_α if either ψ is 'close' to D_α in the sense of RD (i.e., $R(\psi\|D_\alpha)$ is 'small'), or (as below) if ψ is sufficiently 'wider' than a D_α so that $R(\psi\|\psi + D_\alpha)$ is 'small'.

Comparison with other reductions. The following table is providing a comparison between this result and the previous ones mentioned in Section 2.3. The advantage of our result is to have a smaller growth between α and β compared to the first one and to be dimension preserving compared to the second one with only a loss in $\sqrt{\log n}$ between α and β.

Param.	[DM13]	using [BGM+16]	Th. 3.1
m	fixed	-	-
n'	$n/2$	$n/\log q$	n
β	$\geq m\alpha \cdot n^{\sigma}$	$\Omega(m\alpha/\sqrt{\log n})$	$\Omega(m\alpha/\log n)$

Table 3.1: Comparing the main parameters of different reductions from $\mathrm{LWE}_{n',q,D_\alpha}$ to $\mathrm{LWE}_{n,q,U_\beta}$ for a fixed n and another parameter $\sigma \in (0,1)$.

3.1.3 Alternative reduction from LWE to LWR

Another application of the Rényi divergence is to provide an alternative proof that the Learning With Rounding (see Section 2.5) problem [BPR12] is no easier than LWE. This reduction is published in [BLR+18] and is described more in details in Section 3.3.1.

3.2 New results on the hardness of Module LWE

I proposed this direction of work to one of my PhD student, Katharina Boudgoust, and to a post-doc working with me, Weiqiang Wen. Corentin Jeudy joined us on this research during his first internship with us. As explained previously, the hardness of the binary secret version of the Ring LWE problem is an important and difficult open problem. As several reductions exists to show the hardness of the binary version of LWE, our first goal was to apply them to the module variant, and see for which rank of the module they could be adapted. The results detailed in this section are from two research articles, both with Katharina, Corentin and Weiqiang. The first one [BJRW20] is published in the proceeding of Asiacrypt 2020, and the second one [BJRW21] is to appear in the proceeding of CT-RSA 2021.

3.2.1 Hardness of MLWE with binary secret

Our first main contribution in [BJRW20] is a reduction from MLWE to bin-MLWE, its binary secret version, if the module rank d is at least of size $\log_2 q + \omega(\log_2 n)$, where n denotes the degree of the underlying number field and q the modulus. To the best of our knowledge, this is the first result on the hardness of a structured variant of LWE with binary secret.

Our proof follows the proof structure of Goldwasser et al. [GKPV10], but achieves better parameters by using the Rényi divergence (see Section 3.1), while being as direct and short as the original proof. The improvement on the noise rate $\frac{\alpha}{\beta}$ compared to [GKPV10] stems from the fact that the Rényi divergence only needs to be constant for the reduction to work and not necessarily negligibly close to 1 (compared to negligibly close to 0 for the statistical distance). However, using the Rényi divergence as a tool for distance measurement requires to move to the search variants of MLWE and its binary version, respectively.

Theorem 3.2. *Let K be a cyclotomic number field of degree n with R its ring of integers. Let ℓ, d, m and q be positive integers with q prime and m polynomial in n. Further, let α and β be positive real numbers such that $\frac{\alpha}{\beta} \leq \frac{1}{\sqrt{m \cdot n^2 d}}$. Let ε be a positive real number with $\varepsilon \in [0,1)$ such that $\beta q \geq \eta_\varepsilon(R^\vee)$. and $\varepsilon = O(\frac{1}{m})$.*

Then, for any $d \geq \ell \cdot \log_2 q + \omega(\log_2 n)$*, there is a probabilistic polynomial-time reduction from* search $\mathrm{MLWE}^m_{n,\ell,q,\mathcal{D}_{R^\vee,\beta q}}$ *and* $\mathrm{MLWE}^{m,d}_{n,\ell,q,\mathcal{D}_{R^\vee,\alpha q}}$ *to search* bin-$\mathrm{MLWE}^m_{n,d,q,\mathcal{D}_{R^\vee,\beta q}}$.

The degree n of the number field K and the number of samples m are preserved. The reduction increases the rank of the module from ℓ to $\ell \cdot \log_2 q + \omega(\log_2 n)$ and the Gaussian width from αq to $\alpha q \cdot \sqrt{m} \cdot n^2 d$. Further, $\mathrm{MLWE}^m_{n,\ell,q,\mathcal{D}_{R^\vee,\alpha q}}$ trivially reduces to search $\mathrm{MLWE}^m_{n,\ell,q,\mathcal{D}_{R^\vee,\beta q}}$, as $\beta \geq \alpha$.

Within the proof of Theorem 3.2 we need to apply a leftover hash lemma over rings (detailed in our article [BJRW20, Lemma 7] and adapted from [Mic07]), and thus it requires the modulus q to be prime. Furthermore, we also need to use cyclotomic number fields $K = \mathbb{Q}(\zeta)$, where ζ is a primitive root of unity, as in this case we have a direct relation between $R^\vee$ and R as explained in Section 2.4.

Related work. We now compare the results of Theorem 3.2 with the former results on LWE (described in Section 2.3.2). The LWE problem can be seen as a special case of Module LWE, where the ring is $\mathbb{Z}$ and the degree n equals 1. In this case, the rank ℓ of the module corresponds to the dimension of the LWE problem and should be polynomial in the security parameter. Hence, the error-ratio is given by $\beta \geq \alpha\sqrt{m} \cdot d$ and $d \geq \ell \log_2 q + \omega(\log_2 \ell)$. Asymptotically, we lose a factor of $\sqrt{d}$ in the error-ratio in our reduction compared to the former results for LWE [BLP$^+$13, Mic18]. However, our proof is as direct and short as the original one in [GKPV10]. We don't need to define intermediate problems such as *First-is-errorless* LWE and *Extended* LWE as in [BLP$^+$13] and no gadget matrix construction as in [Mic18]. Note that adapting the proof of [Mic18] asks to define a corresponding gadget matrix, which does not seem to work in an obvious way. By replacing the statistical distance by the Rényi divergence and switching to the search variants we obtain a much better result than in the original paper from Goldwasser et al. [GKPV10].

Alternative reduction. In a second paper [BJRW21], we propose an alternative approach to prove the hardness of Module LWE with binary secrets over cyclotomic fields. This reduction from MLWE to bin-MLWE also holds if the module rank d of the binary version is at least of size $\log_2 q + \omega(\log_2 n)$, where n denotes the degree of the underlying number field and q the modulus. However, our noise growth is smaller as our Gaussian parameter only increases by a factor $n\sqrt{2d}\sqrt{4n^2+1} = \Theta(n^2\sqrt{d})$ for cyclotomics. This new reduction removes the dependency in m in the noise ratio $n^2 d\sqrt{m}$ present in Theorem 3.2, which is more advantageous as we usually take $m = O(n \log_2 n)$ samples, and also gains an extra factor $\sqrt{d}$. Additionally, when bridging to LWE, the noise ratio is improved to $\sqrt{10d}$, our work thus matches the results from Brakerski et al. [BLP$^+$13] when we take the ring R to be of degree 1.

Theorem 3.3 (Informal)**.** *For a cyclotomic field of degree* n*, the* bin-MLWE *problem with rank* d *and Gaussian parameter less than* β *is at least as hard as* MLWE *with rank* k *and Gaussian parameter* α*, if* $d \geq (k+1)\log_2 q + \omega(\log_2 n)$ *and* $\beta/\alpha = \Theta(n^2\sqrt{d})$*, where* q *is a modulus such that the cyclotomic polynomial has a specific splitting behavior in* $\mathbb{Z}_q[x]$*.*

The condition on q can be relaxed by switching the modulus from q to any polynomially large p using [LS15, Thm. 4.8]. This entails an additional increase in the Gaussian noise by a factor of $\max(1, q/p) \cdot n^{3/4} d^{1/2} \cdot \omega((\log_2 nd)^2)$. The overall noise ratio then increases to $\beta/\alpha = \max(1, q/p) \cdot n^{11/4} d \cdot \omega((\log_2 nd)^2)$.

In the hope of achieving better parameters than [BJRW20], which is inspired by the proof of [GKPV10], we follow the proof idea of Brakerski et al. [BLP$^+$13] by introducing

the two intermediate problems first-is-errorless MLWE and Extended-MLWE. We first reduce MLWE to the first-is-errorless MLWE variant, where the first sample is not perturbed by an error. We then reduce the latter to Extended-MLWE, which can be seen as MLWE with an extra information on the error vector $\mathbf{e}$ given by $\langle \mathbf{e}, \mathbf{z} \rangle$ for a uniformly chosen $\mathbf{z}$ in the set of binary ring elements set $\mathcal{Z} = (R_2^\vee)^d$. We then use a lossy argument, relying on the newly derived Extended-MLWE hardness assumption and a ring version of the leftover hash lemma, to reduce Extended-MLWE to bin-MLWE. An overview of the full reduction is provided in Figure 3.1.

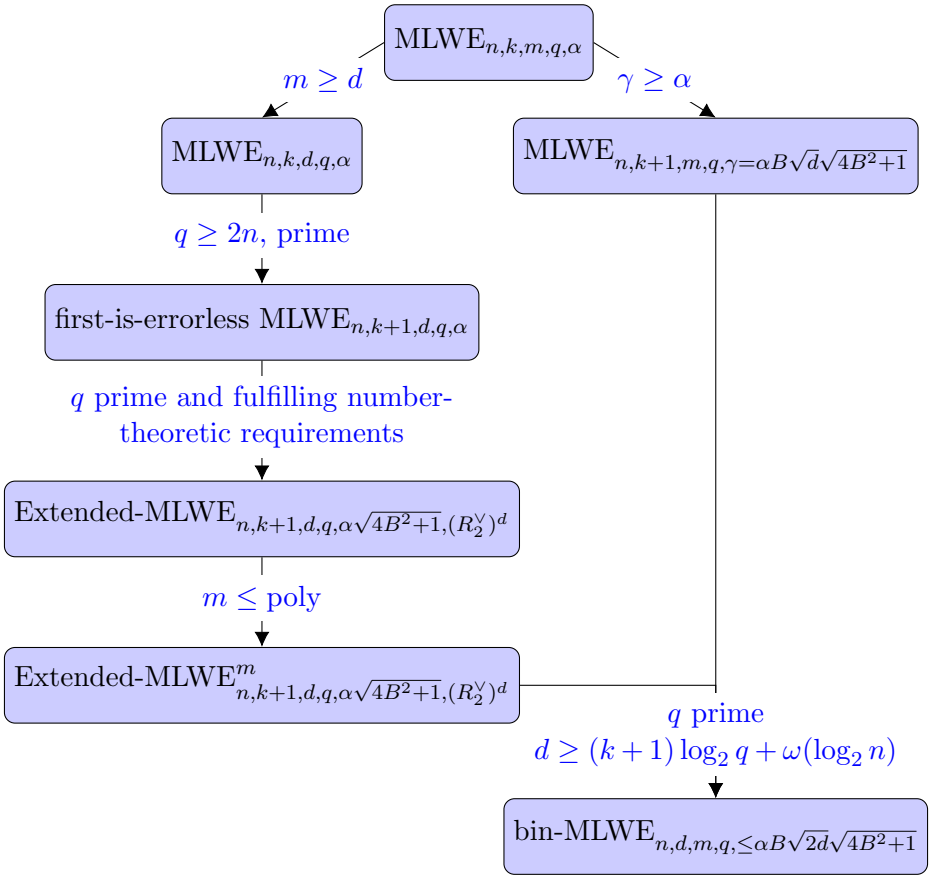

Figure 3.1: Summary of the formal proof of Theorem 3.3, where $B = \max_{x \in R_2} \|\sigma(x)\|_\infty$ and σ is the canonical embedding. In cyclotomic fields, we have $B \leq n$. The assumptions on q concern the splitting behavior of the cyclotomic polynomial in $\mathbb{Z}_q[x]$, and are discussed in the paper.

The main challenge of this reduction is the use of matrices composed of ring elements. The proof in [BLP+13, Sec. 4] requires the construction of unimodular matrices which is not straightforward to adapt in the module setting because of invertibility issues. Our construction relies on units of the quotient ring R/qR, which are much harder to describe than the units of $\mathbb{Z}/q\mathbb{Z}$ to say the least. This is the reason why we need to control the splitting structure of the cyclotomic polynomial modulo q. This requires q to satisfy certain number-theoretic properties and to be sufficiently large so that all the non-zero binary ring elements are units of R_q. The second complication comes from using both the coefficient embedding and the canonical embedding (we refer to [BJRW21, Section 2.1] for the formal definition of those two embeddings). Even though some manipulations on

Gaussian distributions require the use of the canonical embedding, we choose the secret to be binary in the coefficient embedding rather than the canonical embedding. For power-of-two cyclotomics, using the canonical embedding for binary secrets requires the rank d to be larger by a factor n than when using the coefficient embedding.

3.2.2 Classical hardness of M-LWE for a linear rank

Our second main contribution is to use the result from Section 3.2.1 to prove a classical reduction from Approx-GapSVP over Modules, with module rank at least 2, to MLWE for any polynomial-sized modulus $\hat{p}$ and module rank d at least $2n + \omega(\log_2 n)$, for the case of power-of-two cyclotomics.

At a high level, we follow the structure of the classical hardness proof of LWE from [BLP+13]. Overall, we need three ingredients:

- First, the classical hardness of Module LWE with an exponential-sized modulus, where we adapt the proof of Peikert [Pei09] using adapted results from [PRS17],
- As a second component, we need the hardness of Module LWE using a binary secret (using Theorem 3.2 explained in Section 3.2.1),
- Finally, a modulus reduction technique, adapted from [AD17], this technique is using the binary secret variant of Module LWE to provide a reduction from MLWE with an exponential modulus to MLWE with a polynomial modulus.

More formally, achieving this full proof requires several technical steps that need to be carefully handled. Figure 3.2 gives an overview of the full proof as provided in [BJRW20].

Theorem 3.4. *Let ν be a power of 2, defining the ν-th cyclotomic number field with R its ring of integers of degree $n = \nu/2$. Let $d, \hat{p}, m$ be positive integers and $\hat{\beta}$ and γ be positive reals. Fix $\varepsilon \in (0, \frac{1}{2})$ such that $\hat{\beta} \geq \sqrt{2} \cdot 2^{-n} \cdot \eta_\varepsilon(R^\vee)$ and $\varepsilon = O\left(\frac{1}{m}\right)$. There is a* classical *probabilistic polynomial-time reduction from* Approx-GapSVP *over Modules$_\gamma$ to* $\mathrm{MLWE}^m_{n,d,\hat{p},\Upsilon_{\hat{\beta}}}$*, where $d \geq 2n + \omega(\log_2 n)$ and*

$$\hat{\beta} = \tilde{\Theta}\left(\frac{\sqrt{m} \cdot n^{\frac{21}{4}}}{\gamma}\right).$$

Classical hardness of Ring LWE. The first result about the classical hardness of Ring LWE with exponential-sized modulus has been informally mentioned in [BLP+13]. It can be achieved in two steps. First, by a dimension-modulus switching as in [BLP+13], LWE in dimension d and modulus q can be reduced to LWE in dimension 1 and modulus q^d with a slightly increased error rate. Then, by a ring switching technique as in [GHPS12], the latter one can be reduced to Ring LWE over a ring of any degree n and modulus q^d, while keeping the same error rate. For more details on the second step, we refer to [AD17, App. B].

On the other hand, as a direct application of our classical hardness result of Module LWE, we can provide an alternative solution for the classical hardness result of Ring LWE with exponential-sized modulus. The idea is that, using a rank-modulus switching as in [WW19], we can instead reduce from MLWE over d-rank modules of n-degree ring and modulus q, to RLWE with n-degree ring and modulus q^d, with a slightly increased error rate. However, we remark that the underlying worst-case lattice problems are different for these two results. Suppose that we consider the classical hardness of Ring LWE

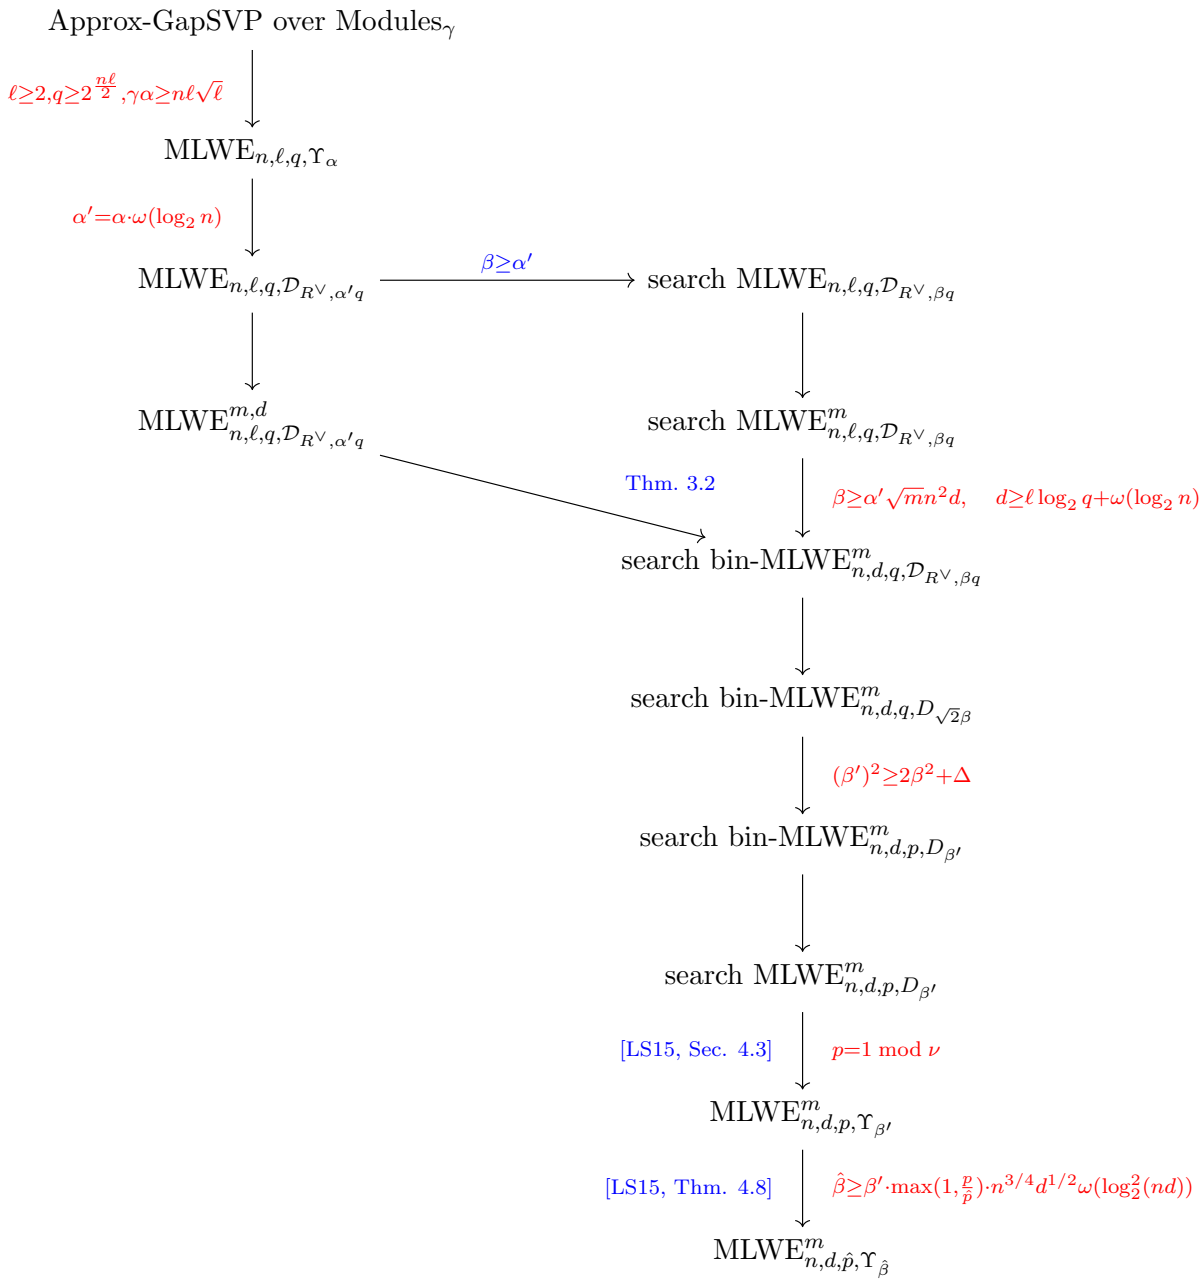

Figure 3.2: Overview of the complete classical hardness proof of MLWE for linear rank d and arbitrary polynomially large modulus $\hat{p}$, as stated in Theorem 3.4 for K the ν-th cyclotomic number field of degree n. The parameter Δ is determined by the underlying ring R and is poly(n) for the case of power-of-two cyclotomics.

over n-degree ring and q^d modulus where $d = \mathcal{O}(n)$. Then, the underlying problem is the standard GapSVP over general lattices of dimension $\mathcal{O}(\sqrt{n})$ for the first result, while it is Approx-GapSVP over Modules over rank-2 modules of $\mathcal{O}(n)$-degree ring for the second one.

3.3 New results on the Learning With Rounding problem

We also investigated the hardness of the Learning With Rounding problem, which is a deterministic variant of LWE.

3.3.1 Alternative reduction from LWE to LWR

Another application of the Rényi divergence defined in Section 3.1 is to provide an alternative proof that the Learning With Rounding (see Section 2.5) problem [BPR12] is no easier than LWE. This reduction is published in [BLR$^+$18], the journal version of [BLL$^+$15].

Our reduction is the first which preserves the dimension n without resorting to noise flooding (which significantly degrades the noise rate): the reductions from [AKPW13, BGM$^+$16] do not preserve the dimension, and the one from [BPR12] preserves the dimension but makes use of noise flooding. In [AKPW13], the authors can get close to preserve the dimension up to a constant but at a price of larger polynomial moduli. Denoting by $\mathbb{Z}_p$ the ring in which we perform rounding, our new reduction also gains extra factors of $p\sqrt{\log n}$ and $pn\sqrt{\log n}$ in the number of LWR samples handled, compared to [BGM$^+$16] and [AKPW13], respectively.

We first combine (in Theorem 3.5) the hardness result of [BGM$^+$16] for LWR with other results, to state it as a reduction from the standard LWE problem, so that it would be comparable with our alternative reduction. This result of [BGM$^+$16] makes use of the Rényi divergence and was inspired by an earlier version of our work [BLL$^+$15], within a proof that can be seen as a variant of the [MM11] search-to-decision reduction for LWE.

Theorem 3.5 (Using [BGM$^+$16, Theorem 3]). *Let $qm = O(poly(n))$, and $n \leq m \leq O\left(\frac{\sqrt{\log n}}{p\alpha}\right)$. Then there is a polynomial-time reduction from $\mathrm{LWE}_{n/\log q,q,D_\alpha,m}$ to $\mathrm{LWR}_{n,q,p,m}$.*

This reduction can be obtained in the following five steps:

- A reduction from decisional $\mathrm{LWE}_{n/\log q,q,D_\alpha,m}$ to decisional bin-$\mathrm{LWE}_{n,q,D_\alpha,m}$, from LWE to its binary variant, using [BLP$^+$13].
- A trivial reduction from decisional bin-$\mathrm{LWE}_{n,q,D_\alpha,m}$ to search bin-$\mathrm{LWE}_{n,q,D_\alpha,m}$,
- A reduction from search bin-$\mathrm{LWE}_{n,q,D_\alpha,m}$ to search bin-$\mathrm{LWE}_{n,q,D'_{\alpha,B'},m}$, with $D'_{\alpha,B'}$ the distribution D_α truncated (by rejection) to the interval $[-B', B']$, using the Rényi divergence.
- A reduction from search bin-$\mathrm{LWE}_{n,q,D'_{\alpha,B'},m}$ to search bin-$\mathrm{LWE}_{n,q,\phi,m}$, with $\phi = \frac{1}{q}\lfloor qD'_{\alpha,B'}\rceil$, from a continuous to a discrete error distribution.
- A reduction from search bin-$\mathrm{LWE}_{n,q,\phi,m}$ to decisional $\mathrm{LWR}_{n,q,p,m}$ via [BGM$^+$16, Theorem 3].

We now give a tighter reduction than above from LWE to LWR for composite q.

Theorem 3.6 (Adapted from [BGM$^+$16, Th. 13]). *Let p and q be two integers such that p divides q, $\beta = q/(2p)$ and $m' = m \cdot q/p$. There is a polynomial time reduction from $\mathrm{LWE}_{n,q,U_\beta,m'}$ to $\mathrm{LWR}_{n,q,p,m}$.*

Then, we show (in Theorem 3.7) a new dimension-preserving hardness result for LWR, obtained by composing the previous hardness result for LWE with uniform noise with another reduction from [BGM+16] that reduces LWE with uniform noise to LWR. Interestingly, our new reduction for LWR also makes use of the Micciancio-Mol reduction [MM11], but unlike the LWR reduction in [BGM+16, Theorem 3], ours uses [MM11] as a black box within the reduction of Theorem 3.1.

Theorem 3.7. *Let p divide q, $m' = m \cdot q/p$ with $m = O(\log n/\alpha)$ for $m' \geq m \geq n \geq 1$. There is a polynomial-time reduction from $\mathrm{LWE}_{n,q,D_\alpha,m'}$ to $\mathrm{LWR}_{n,q,p,m}$.*

Let $\beta = q/(2p)$. The reduction has two steps:

- A reduction from $\mathrm{LWE}_{n,q,D_\alpha,m'}$ to $\mathrm{LWE}_{n,q,U_\beta,m'}$, using Theorem 3.1.
- A reduction from $\mathrm{LWE}_{n,q,U_\beta,m'}$ to $\mathrm{LWR}_{n,q,p,m}$, using Theorem 3.6.

Comparison with previous works. Table 3.2 compares the parameters of Theorems 3.5 from [BGM+16] and 3.7 (our result), and a reduction from [AKPW13]. The reduction in Theorem 3.5 loses a $\log q$ factor in dimension, while our uniform-noise reduction preserves the dimension, which is the first of its kind without resorting to the noise-flooding technique (as [BPR12]). On the downside, our reduction does not preserve the number of samples.

Param.	[AKPW13, Th. 4.1]	Th. 3.5 ([BGM+16])	Th. 3.7
n'	$O\left(\frac{n\log(2\tau)}{\log q}\right)$	$\frac{n}{\log q}$	n
m	$O\left(\frac{\sqrt{\log n}}{\tau n p \alpha}\right)$	$O\left(\frac{\sqrt{\log n}}{p\alpha}\right)$	$O\left(\frac{\log n}{\alpha}\right)$
m'	m	m	$m \cdot \frac{q}{p}$

Table 3.2: Comparing the main parameters of different reductions from $\mathrm{LWE}_{n',q,D_\alpha,m'}$ to $\mathrm{LWR}_{n,q,p,m}$ for a fixed n and another flexible parameter $\tau \geq 1$.

Note that setting $\tau = 1$ gives n' equal to that of Theorem 3.5, while it loses an extra factor n in the denominator of m. On the other hand, setting $\tau = q$ allows for approximately $n = n'$, however for an expense of much smaller m. The reduction in Theorem 3.5 also restricts the number of LWR samples m by a further $O\left(p\sqrt{\log n}\right)$ factor in comparison to our results. This factor is equal to $O\left(\tau p n\sqrt{\log n}\right)$ if we compare our result with that of Theorem 4.1 from [AKPW13].

3.3.2 Middle-Product Learning With Rounding

We started working on the Middle-Product variant of LWE (defined in Section 2.4) at the beginning of Katharina's PhD with Weiqiang Wen, who was also starting a post-doc. Our first goal was to study the possibility of building a signature scheme proven secure under the MP-LWE assumption, and ideally by building a trapdoor. But, discussing with Shi Bai, who was working on the same topic with Dipayan Das and Zhenfei Zhang, we decided to work together on adapting the Middle Product assumption to the Learning With Rounding problem (defined in Section 2.5). The following result is from the article [BBD+19], published in the conference Asiacrypt 2019. Note that this work was published before the recent result of [LW20] on the hardness of LWR over rings.

Our first main contribution is a new hardness assumption which we refer to as the *Middle Product Computational Learning With Rounding* (MP-CLWR) problem. On one hand, the MP-CLWR problem uses rounding in a similar way to LWR and hence avoids the error sampling. On the other hand, the MP-CLWR problem is analogue to the MP-LWE problem whose hardness does not depend on a specific polynomial. Thus, the MP-CLWR assumption enjoys the desired properties from both, the security advantage of MP-LWE and the simplicity advantage of LWR. We show that the MP-CLWR problem is at least as hard as the decisional MP-LWE problem parametrized over a set of polynomials. To complete the reduction, we also bring in some new results on random Hankel matrices which might be of independent interest. As a typical application, we propose a PKE scheme based on this MP-CLWR assumption which is IND-CPA secure in the random oracle model. The attractiveness of our encryption scheme stems from the fact that we only have to round the middle-product of two polynomials instead of sampling Gaussian error during public key generation while guaranteeing the same security and having the same asymptotic key and ciphertext sizes as [RSSS17].

The MP-CLWR problem. An MP-CLWR sample is given by $(a, b = \lfloor a \odot_d s \rceil_p)$, where a is sampled from the uniform distribution over $\mathbb{Z}_q^{<n}[x]$ and s is a fixed element in $\mathbb{Z}_q^{<n+d-1}[x]$. We define the MP-CLWR problem as the following game, where we embed the MP-CLWR samples into two experiments. In both experiments, three different parties appear: A challenger $\mathcal{C}$, an adversary $\mathcal{A}$ and a source $\mathcal{S}$. The source $\mathcal{S}_1$ of the first experiment provides t different MP-CLWR samples $(a_i, \lfloor a_i \odot_d s \rceil_p)_{i\in[t]}$ and the source $\mathcal{S}_2$ of the second experiment provides t rounded uniform samples $(a_i, \lfloor b_i \rceil_p)_{i\in[t]}$, where all a_i and b_i are independently sampled from the corresponding uniform distribution. The challenger $\mathcal{C}$ now uses these samples to compute an Input and a Target. It then sends the Input to the adversary $\mathcal{A}$ which itself computes an Output. The adversary wins the experiment if Target = Output. The important point in this setting is that the challenger $\mathcal{C}$ and the adversary $\mathcal{A}$ are the same in both experiments. The MP-CLWR assumption captures the fact that an adversary has no more advantage to compute the correct output if it receives rounded middle-product samples than if it gets rounded uniform samples. A formal definition of MP-CLWR is given in [BBD+19, Section 4.1].

Our reduction from MP-LWE to MP-CLWR is dimension-preserving and works for polynomial-sized modulus q. The parameters d and n describe the order of the middle-product, t denotes the number of samples and p defines the rounding.

Theorem 3.8 (Hardness of MP-CLWR)**.** *Let d, n, p, q and t be positive integers with $0 < d \leq n$ and $q \geq p \geq 2$. Further, let $q = \prod_{i\in[l]} p_i^{\alpha_i}$ be the prime power factorization of q with some $l > 0$, where p_i is prime and $\alpha_i > 0$ for all $i \in [l]$. Let χ be an error distribution over $\mathbb{R}^{<d}[x]$ which is balanced and B-bounded with probability at least δ, fulfilling $q > 2pBdt$ and $\delta \geq 1 - \frac{1}{td}$. There is a reduction from the decisional* MP-LWE$_{q,n,d,\chi}$ *problem to the* MP-CLWR$_{p,q,n,d,t}$ *problem, with t the number of samples fixed beforehand.*

Assume that the advantage of an MP-CLWR *solver is ε. Then, there is an* MP-LWE *solver with advantage at least*

$$\left(\frac{1}{e^2}\left(\varepsilon + \mathcal{Q}_{\mathcal{C},\mathcal{A}}\right)^2\right) \cdot \prod_{i\in[l]} \left(1 - \frac{1}{p_i}\right).$$

In order to prove the theorem, we show the following sequence of reductions where the lemma are proven in the full version of the paper:

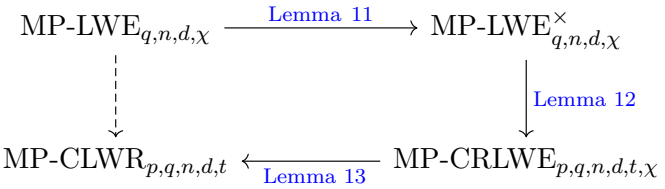

The first part of this sequence, [BBD$^+$19, Lemma 11], gives a reduction from decisional MP-LWE to decisional MP-LWE$^{\times}$, where the latter one denotes the MP-LWE problem where the secret is sampled uniformly at random from the set of elements having full rank Hankel matrix. The Hankel matrix plays an important role during the reductions as one can use it to represent the middle-product. We give a lower bound of the probability that the Hankel matrix of a random element has full rank and prove a uniformity property of the middle-product. This property is used in Lemma [BBD$^+$19, Lemma 12], where we show a reduction from the rounded middle-product LWE problem to the middle-product LWR problem, for their computational versions. Note that using the Rényi divergence requires the requested number of samples t to be fixed a priori. This is a necessary requirement which is also imposed in [BGM$^+$16] and [CZZ18].

Chapter 4

Conclusion

As explained in the introduction, lattice-based cryptography is a promising post-quantum alternative to modern asymmetric cryptography. One important challenge today is to be prepared for this change, as many companies are working on building quantum computers. Since my PhD thesis in 2014, the NIST competition has changed the landscape of this approach by encouraging a lot the research on practical public key encryption and signature schemes. Nevertheless, there are still many open questions to ensure a transition to a post-quantum cryptography, with all its applications, both in theory and practical aspects.

My research takes place in this context. The works I presented more in details focus on studying the theoretical hardness of the Learning With Errors problem and its variants, both by improving existing reductions and by building new ones. In my opinion, reductions are very important as a security foundation of lattice-based cryptography, even if in practice, the parameters used for cryptographic constructions are not necessarily the ones induced by the security reduction. Another limitation is that most of the reductions are not tight, in the sense that in practice, if an adversary manages to attack a "proven secure" scheme, it would not really give an *efficient* algorithm to solve the worst-case hard problem on lattice. As the details on the relation between the advantages or the precise time of the reduction are often not given, this may be complicated to study for practical applications.

But it is not necessarily the only use we have today for all those reductions. Indeed, we all know that at some point, cryptography always relies on a conjecture. The important question is to know how confident we are that this conjecture is true. From my point of view, reductions are a very useful tool to increase this confidence in the security of the schemes. They allow a better understanding of the hardness of the problems, and of the links between all of them. It seems particularly important to me to better understand the structured variants, as they will probably be used in the next standardized signatures or public key encryptions, and the variants using other error or secret distributions, as they seem also to be the ones we want to use in practice.

The second approach to gain in confidence is of course cryptanalysis, which studies the security by attacking directly a construction or a protocol, or by finding a better algorithm to solve the hard problems on which it depends. Then in practice, we choose concrete parameters for a cryptographic construction depending on the best known attack. The parameters obtained for a given constructions from the two approaches are still quite different. If we stick to the theoretical reductions to choose them, the scheme will be quite inefficient in practice. This is a reason why it is still interesting to improve the existing reductions as it could allow to get closer from parameters obtained using the proven security instead of parameters obtained from the best-known attacks.

I really enjoy working on this research topic and I will continue in the following directions, which I hope will help for this transition to post-quantum cryptography.

Specific opens problems. I start with the open problems specific to each section in Chapter 3, except for the first one concerning the Rényi divergence as this work is different from the others, by introducing a tool to obtain better proofs instead of providing new ones.

- **On the hardness of Module LWE with binary secret.**
  - First reduction in [BJRW20]: In this result, we incurred several restrictions on the class of number fields we look at. The hardness proof of search bin-MLWE (Theorem 3.2) is restricted to cyclotomic fields, in order to bound the norm of the Vandermonde matrix in the proof.
  - Alternative reduction in [BJRW21]: Most of our results rely on the class of number fields $K = \mathbb{Q}(\zeta)$ where the ring of integers is $R = \mathbb{Z}[\zeta]$. Although this class includes all cyclotomic fields, we leave as an open problem to generalize these results to a larger class of number fields.

    The construction of the matrix in [BJRW21, Lemma 15] seems optimized in terms of its impact on the Gaussian parameter. However, its invertibility restricts the underlying number field, as well as the structure of the chosen modulus q. A better understanding of the unit group of R_q for general cyclotomic fields and other number fields might help relax the restrictions on the modulus q for the reduction to go through.

  For both results, the hardness of search bin-MLWE for a rank which is smaller than $\log_2 q + \omega(\log_2 n)$, in particular for binary RLWE, is still an open problem. In practice, we usually chose a small constant rank (< 10), as for instance in the submission to the NIST standardization process Kyber [BDK+18].

- **On the classical hardness of LWE.**

  For this result also, the classical hardness of MLWE for a rank which is smaller than $\log_2 q + \omega(\log_2 n)$, in particular for RLWE with a polynomial modulus is still an open problem.

  Further, quantifying the error increase in the modulus switching reduction for other number fields than power-of-two cyclotomics may be interesting. The current bounds heavily depend on the singular values of the secret's rotation matrix, which further depend on the underlying number field.

- **On the MP-CLWR problem.**

  The open problems mentioned in the published article, i.e. a reduction from decisional RLWE to decisional RLWR with a polynomial-sized modulus, was later given in [LW20], but using a different rounding operation. In the middle-product setting, it would also be of interest to show a reduction from decisional MP-LWE to decisional MP-LWR. Such a hardness result would help to build a secure encryption scheme based on the decisional MP-LWR in the standard model.

Hardness of Learning With Errors and its variants. To conclude on more generals open problems on the hardness of the Learning With Errors problem and its variants, I think the actual hardness of Ring LWE is one of the major open question. The second one is on having a better understanding of the hardness of Module LWE, which seems to be mostly used in practice, and more generally of all the variants used in the security of NIST finalists constructions.

Concerning the hardness of Ring LWE, what we know today is that there exists a quantum reduction from Approx-SVP on ideal lattices to Ring LWE. A first question would be to know if there exist a classical reduction, as for LWE, from ideal lattice to Ring LWE. Indeed the results obtained for LWE, or Module LWE in the linear rank case, does not adapt for the Ring setting. An intermediate step to show this result could be to prove the hardness of Ring LWE with a binary secret, which is still an open question. But more importantly, many recent works [CGS14, EHKS14, BS16, CDPR16, CDW17, DPW19, PHS19], followed by our work in [BR20] with Olivier Bernard, suggest that the Approx-SVP problem on ideal lattices could not be as hard as Approx-SVP in arbitrary lattices, in particular in the quantum setting. Note that if it appears that Approx-SVP on ideal lattices is not quantumly hard, it would not mean that it is also the case of Ring LWE as the reduction only shows that Ring LWE is at least as hard as Approx-SVP on ideal lattices, and as there is no converse reduction. Works on Ring and Module LWE show that Ring-LWE is equivalent to solving Approx-SVP in a module of rank 2. Then the gap in difficulty of Module Approx-SVP seems to be between ideals lattices (rank 1) and module lattices on rank 2. An interesting question would be to understand better this gap, and determine if Approx-SVP on ideal is really easier than Approx-SVP on modules of rank 2 and why.

The module setting could be a solution, as it seems to be a nice trade-off between security and efficiency. But there are still open problems to study, as if it is possible to obtain better parameters than a linear rank case for our classical reduction. The case of a constant rank would be particularly interesting, as it is the case used in NIST submissions. We also saw that NIST submissions are considering other distributions of the secret or the error. For instance Kyber [BDK$^+$18] has its security relying on Module LWE with a binomial error and secret. Obtaining a reduction to show the hardness of this particular variant, or even of LWE with both error and secret binomial, would help to evaluate the security of this scheme. An other example is concerning Fully Homomorphic encryption, where the secret distribution is binary, and even with a constant or bounded number of ones smaller than one half.

Finally, as discussed above, an important open problem for cryptography would be to obtain a full serie of reductions with good enough parameters to be used in practice.

Bibliography

[AA16] Jacob Alperin-Sheriff and Daniel Apon. Dimension-preserving reductions from LWE to LWR. *IACR Cryptol. ePrint Arch.*, 2016:589, 2016.

[ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*, pages 98–115. Springer, 2010.

[ACLL15] Martin R. Albrecht, Catalin Cocis, Fabien Laguillaumie, and Adeline Langlois. Implementing candidate graded encoding schemes from ideal lattices. In *ASIACRYPT (2)*, volume 9453 of *Lecture Notes in Computer Science*, pages 752–775. Springer, 2015.

[ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618. Springer, 2009.

[AD17] Martin R. Albrecht and Amit Deo. Large modulus ring-lwe $\geq$ module-lwe. In *ASIACRYPT (1)*, volume 10624 of *Lecture Notes in Computer Science*, pages 267–296. Springer, 2017.

[AG11] Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In *ICALP (1)*, volume 6755 of *Lecture Notes in Computer Science*, pages 403–415. Springer, 2011.

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108. ACM, 1996.

[Ajt98] Miklós Ajtai. The shortest vector problem in L_2 is *NP*-hard for randomized reductions (extended abstract). In *STOC*, pages 10–19. ACM, 1998.

[AKPW13] Joël Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. Learning with rounding, revisited - new reduction, properties and applications. In *CRYPTO (1)*, volume 8042 of *Lecture Notes in Computer Science*, pages 57–74. Springer, 2013.

[AKS01] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610. ACM, 2001.

[AR05] Dorit Aharonov and Oded Regev. Lattice problems in NP cap conp. *J. ACM*, 52(5):749–765, 2005.

[BBD+19] Shi Bai, Katharina Boudgoust, Dipayan Das, Adeline Roux-Langlois, Weiqiang Wen, and Zhenfei Zhang. Middle-product learning with rounding problem and its applications. In *ASIACRYPT (1)*, volume 11921 of *Lecture Notes in Computer Science*, pages 55–81. Springer, 2019.

[BBF+19] Hayo Baan, Sauvik Bhattacharya, Scott R. Fluhrer, Óscar García-Morchón, Thijs Laarhoven, Ronald Rietman, Markku-Juhani O. Saarinen, Ludo Tolhuizen, and Zhenfei Zhang. Round5: Compact and fast post-quantum public-key encryption. In *PQCrypto*, volume 11505 of *Lecture Notes in Computer Science*, pages 83–102. Springer, 2019.

[BCE+20] Samuel Bouaziz-Ermann, Sébastien Canard, Gautier Eberhart, Guillaume Kaim, Adeline Roux-Langlois, and Jacques Traoré. Lattice-based (partially) blind signature without restart. *IACR Cryptol. ePrint Arch.*, 2020:260, 2020.

[BD20] Zvika Brakerski and Nico Döttling. Hardness of LWE on general entropic distributions. In *EUROCRYPT (2)*, volume 12106 of *Lecture Notes in Computer Science*, pages 551–575. Springer, 2020.

[BDK+18] Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - kyber: A cca-secure module-lattice-based KEM. In *EuroS&P*, pages 353–367. IEEE, 2018.

[BFRS18] Pauline Bert, Pierre-Alain Fouque, Adeline Roux-Langlois, and Mohamed Sabt. Practical implementation of ring-sis/lwe based signature and IBE. In *PQCrypto*, volume 10786 of *Lecture Notes in Computer Science*, pages 271–291. Springer, 2018.

[BGG+14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *EUROCRYPT*, volume 8441 of *Lecture Notes in Computer Science*, pages 533–556. Springer, 2014.

[BGM+16] Andrej Bogdanov, Siyao Guo, Daniel Masny, Silas Richelson, and Alon Rosen. On the hardness of learning with rounding over small modulus. In *TCC (A1)*, volume 9562 of *Lecture Notes in Computer Science*, pages 209–224. Springer, 2016.

[BGPW16] Johannes Buchmann, Florian Göpfert, Rachel Player, and Thomas Wunderer. On the hardness of LWE with binary error: Revisiting the hybrid lattice-reduction and meet-in-the-middle attack. In *AFRICACRYPT*, volume 9646 of *Lecture Notes in Computer Science*, pages 24–43. Springer, 2016.

[BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, pages 309–325. ACM, 2012.

[BHLY16] Leon Groot Bruinderink, Andreas Hülsing, Tanja Lange, and Yuval Yarom. Flush, gauss, and reload - A cache attack on the BLISS lattice-based signature scheme. In *CHES*, volume 9813 of *Lecture Notes in Computer Science*, pages 323–345. Springer, 2016.

[BJRW20] Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois, and Weiqiang Wen. Towards classical hardness of module-lwe: The linear rank case. In *ASIACRYPT (2)*, volume 12492 of *Lecture Notes in Computer Science*, pages 289–317. Springer, 2020.

[BJRW21] Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois, and Weiqiang Wen. On the hardness of module-lwe with binary secret. In *To appear in CT-RSA*, 2021.

[BLL+15] Shi Bai, Adeline Langlois, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the rényi divergence rather than the statistical distance. In *ASIACRYPT (1)*, volume 9452 of *Lecture Notes in Computer Science*, pages 3–24. Springer, 2015.

[BLP+13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584. ACM, 2013.

[BLR+18] Shi Bai, Tancrède Lepoint, Adeline Roux-Langlois, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the rényi divergence rather than the statistical distance. *Journal of Cryptology*, 31(2):610–640, 2018.

[BPR11] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. *IACR Cryptology ePrint Archive*, 2011:401, 2011.

[BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 719–737. Springer, 2012.

[BR20] Olivier Bernard and Adeline Roux-Langlois. Twisted-phs: Using the product formula to solve approx-svp in ideal lattices. In *ASIACRYPT (2)*, volume 12492 of *Lecture Notes in Computer Science*, pages 349–380. Springer, 2020.

[BS99] Johannes Blömer and Jean-Pierre Seifert. On the complexity of computing short linearly independent vectors and short bases in a lattice. In *STOC*, pages 711–720. ACM, 1999.

[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 *SODA*, pages 893–902. SIAM, 2016.

[CDH+18] Cong Chen, Oussama Danba, Jeffrey Hoffstein, Andreas Hulsing, Joost Rijneveld, John M. Schanck, Peter Schwabe, William Whyte, Zhenfei Zhang, Tsunekazu Saito, Takashi Yamakawa, and Keita Xagawa. NTRU a submission to the nist post-quantum standardization effort, 2018.

[CDPR16] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT (2)*, volume 9666 of *LNCS*, pages 559–585. Springer, 2016.

[CDW17] Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-svp. In *EUROCRYPT (1)*, volume 10210 of *LNCS*, pages 324–348. Springer, 2017.

[CGGI16] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. A homomorphic LWE based e-voting scheme. In *PQCrypto*, volume 9606 of *Lecture Notes in Computer Science*, pages 245–265. Springer, 2016.

[CGK+20] Sébastien Canard, Adela Georgescu, Guillaume Kaim, Adeline Roux-Langlois, and Jacques Traoré. Constant-size lattice-based group signature with forward security in the standard model. In *ProvSec*, volume 12505 of *Lecture Notes in Computer Science*, pages 24–44. Springer, 2020.

[CGS14] Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: A cautionary tale,, 2014.

[Cha82] David Chaum. Blind signatures for untraceable payments. In *CRYPTO*, pages 199–203. Plenum Press, New York, 1982.

[CHL+15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In *EUROCRYPT (1)*, volume 9056 of *Lecture Notes in Computer Science*, pages 3–12. Springer, 2015.

[CIV16] Wouter Castryck, Ilia Iliashenko, and Frederik Vercauteren. Provably weak instances of ring-lwe revisited. In *EUROCRYPT (1)*, volume 9665 of *Lecture Notes in Computer Science*, pages 147–167. Springer, 2016.

[CLT13] Jean-Sébastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In *CRYPTO (1)*, volume 8042 of *Lecture Notes in Computer Science*, pages 476–493. Springer, 2013.

[CMM19] Núria Costa, Ramiro Martínez, and Paz Morillo. Lattice-based proof of a shuffle. In *Financial Cryptography Workshops*, volume 11599 of *Lecture Notes in Computer Science*, pages 330–346. Springer, 2019.

[CRS05] David Chaum, Peter Y. A. Ryan, and Steve A. Schneider. A practical voter-verifiable election scheme. In *ESORICS*, volume 3679 of *Lecture Notes in Computer Science*, pages 118–139. Springer, 2005.

[CZZ18] Long Chen, Zhenfeng Zhang, and Zhenfei Zhang. On the hardness of the computational ring-lwr problem and its applications. In *ASIACRYPT (1)*, volume 11272 of *Lecture Notes in Computer Science*, pages 435–464. Springer, 2018.

[DD12] Léo Ducas and Alain Durmus. Ring-lwe in polynomial rings. In *Public Key Cryptography*, volume 7293 of *Lecture Notes in Computer Science*, pages 34–51. Springer, 2012.

[DDLL13] Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In *CRYPTO (1)*, volume 8042 of *Lecture Notes in Computer Science*, pages 40–56. Springer, 2013.

[DKL+18] Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.

[DKR+19] Jan-Pieter D'Anvers, Angshuman Karmakar, Sujoy Sinha Roy, Frederik Vercauteren, Jose Maria Bermudo Mera, Michiel Van Beirendonck, and Andrea Basso. SABER: M-LWR based KEM, 2019.

[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 *AFRICACRYPT*, volume 10831 of *Lecture Notes in Computer Science*, pages 282–305. Springer, 2018.

[DLP14] Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In *ASIACRYPT (2)*, volume 8874 of *Lecture Notes in Computer Science*, pages 22–41. Springer, 2014.

[DM13] Nico Döttling and Jörn Müller-Quade. Lossy codes and a new variant of the learning-with-errors problem. In *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 18–34. Springer, 2013.

[dPLNS17] Rafaël del Pino, Vadim Lyubashevsky, Gregory Neven, and Gregor Seiler. Practical quantum-safe voting from lattices. In *CCS*, pages 1565–1581. ACM, 2017.

[DPW19] Léo Ducas, Maxime Plançon, and Benjamin Wesolowski. On the shortness of vectors to be found by the ideal-svp quantum algorithm. In *CRYPTO (1)*, volume 11692 of *LNCS*, pages 322–351. Springer, 2019.

[EHKS14] Kirsten Eisenträger, Sean Hallgren, Alexei Y. Kitaev, and Fang Song. A quantum algorithm for computing the unit group of an arbitrary degree number field. In *STOC*, pages 293–302. ACM, 2014.

[FHK+17] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON: Fast-fourier lattice-based compact signatures over NTRU, 2017.

[FOO92] Atsushi Fujioka, Tatsuaki Okamoto, and Kazuo Ohta. A practical secret voting scheme for large scale elections. In *AUSCRYPT*, volume 718 of *Lecture Notes in Computer Science*, pages 244–251. Springer, 1992.

[GGH13a] S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In *Proc. of EUROCRYPT*, volume 7881 of *LNCS*, pages 1–17. Springer, 2013.

[GGH+13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, pages 40–49. IEEE Computer Society, 2013.

[GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In *TCC (2)*, volume 9015 of *Lecture Notes in Computer Science*, pages 498–527. Springer, 2015.

[GHPS12] Craig Gentry, Shai Halevi, Chris Peikert, and Nigel P. Smart. Ring switching in bgv-style homomorphic encryption. In *SCN*, volume 7485 of *Lecture Notes in Computer Science*, pages 19–37. Springer, 2012.

[GKPV10] Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In *ICS*, pages 230–240. Tsinghua University Press, 2010.

[GKV10] S. Dov Gordon, Jonathan Katz, and Vinod Vaikuntanathan. A group signature scheme from lattice assumptions. In *ASIACRYPT*, volume 6477 of *Lecture Notes in Computer Science*, pages 395–412. Springer, 2010.

[GM18] Nicholas Genise and Daniele Micciancio. Faster gaussian sampling for trapdoor lattices with arbitrary modulus. In *EUROCRYPT (1)*, volume 10820 of *Lecture Notes in Computer Science*, pages 174–203. Springer, 2018.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. ACM, 2008.

[GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *STOC*, pages 545–554. ACM, 2013.

[HJ16] Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. In *EUROCRYPT (1)*, volume 9665 of *Lecture Notes in Computer Science*, pages 537–565. Springer, 2016.

[HKLN20] Eduard Hauck, Eike Kiltz, Julian Loss, and Ngoc Khanh Nguyen. Lattice-based blind signatures, revisited. In *CRYPTO (2)*, volume 12171 of *Lecture Notes in Computer Science*, pages 500–529. Springer, 2020.

[HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998.

[Kan83] Ravi Kannan. Improved algorithms for integer programming and related lattice problems. In *STOC*, pages 193–206. ACM, 1983.

[KY19] Shuichi Katsumata and Shota Yamada. Group signatures without NIZK: from lattices in the standard model. In *EUROCRYPT (3)*, volume 11478 of *Lecture Notes in Computer Science*, pages 312–344. Springer, 2019.

[LLL82] Arjen K. Lenstra, Hendrik W. Jr. Lenstra, and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515–534, 1982.

[LLLS13] Fabien Laguillaumie, Adeline Langlois, Benoît Libert, and Damien Stehlé. Lattice-based group signatures with logarithmic signature size. In *ASIACRYPT (2)*, volume 8270 of *Lecture Notes in Computer Science*, pages 41–61. Springer, 2013.

[LLNW14] Adeline Langlois, San Ling, Khoa Nguyen, and Huaxiong Wang. Lattice-based group signature scheme with verifier-local revocation. In *Public Key Cryptography*, volume 8383 of *Lecture Notes in Computer Science*, pages 345–361. Springer, 2014.

[LM06] V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proc. of ICALP (2)*, volume 4052 of *LNCS*, pages 144–155. Springer, 2006.

[LNWX19] San Ling, Khoa Nguyen, Huaxiong Wang, and Yanhong Xu. Forward-secure group signatures from lattices. In *PQCrypto*, volume 11505 of *Lecture Notes in Computer Science*, pages 44–64. Springer, 2019.

[LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *Proc. of EUROCRYPT*, LNCS, pages 1–23. Springer, 2010.

[LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Design Codes and Cryptography*, 75(3):565–599, 2015.

[LSS14] Adeline Langlois, Damien Stehlé, and Ron Steinfeld. Gghlite: More efficient multilinear maps from ideal lattices. In *EUROCRYPT*, volume 8441 of *Lecture Notes in Computer Science*, pages 239–256. Springer, 2014.

[LW20] Feng-Hao Liu and Zhedong Wang. Rounding in the rings. In *CRYPTO (2)*, volume 12171 of *Lecture Notes in Computer Science*, pages 296–326. Springer, 2020.

[Lyu16] Vadim Lyubashevsky. Digital signatures based on the hardness of ideal lattice problems in all rings. In *ASIACRYPT (2)*, volume 10032 of *Lecture Notes in Computer Science*, pages 196–214, 2016.

[Mic98] Daniele Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. In *FOCS*, pages 92–98. IEEE Computer Society, 1998.

[Mic07] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complex.*, 16(4):365–411, 2007.

[Mic18] Daniele Micciancio. On the hardness of learning with errors with binary secrets. *Theory Comput.*, 14(1):1–17, 2018.

[MM11] Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 465–484. Springer, 2011.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718. Springer, 2012.

[MP13] D. Micciancio and C. Peikert. Hardness of SIS and LWE with small parameters. In *Proc. of CRYPTO (1)*, volume 8042 of *LNCS*, pages 21–39. Springer, 2013.

[MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measure. In *Proc. of FOCS*, pages 371–381. IEEE, 2004. Conference version of [MR07].

[MR07] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput*, 37(1):267–302, 2007. Full version of [MR04].

[MV10] Daniele Micciancio and Panagiotis Voulgaris. Faster exponential time algorithms for the shortest vector problem. In *SODA*, pages 1468–1480. SIAM, 2010.

[Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.

[Pei10] Chris Peikert. An efficient and parallel gaussian sampler for lattices. In *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*, pages 80–97. Springer, 2010.

[Pei16] Chris Peikert. How (not) to instantiate ring-lwe. In *SCN*, volume 9841 of *Lecture Notes in Computer Science*, pages 411–430. Springer, 2016.

[Pes16] Peter Pessl. Analyzing the shuffling side-channel countermeasure for lattice-based signatures. In *INDOCRYPT*, volume 10095 of *Lecture Notes in Computer Science*, pages 153–170, 2016.

[PHS19] Alice Pellet-Mary, Guillaume Hanrot, and Damien Stehlé. Approx-svp in ideal lattices with pre-processing. In *EUROCRYPT (2)*, volume 11477 of *LNCS*, pages 685–716. Springer, 2019.

[PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, volume 3876 of *LNCS*, pages 145–166. Springer, 2006.

[PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-lwe for any ring and modulus. In *STOC*, pages 461–473. ACM, 2017.

[Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proc. of STOC*, pages 84–93. ACM, 2005.

[Reg09] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.

[Rén61] Alfréd Rényi. On measures of entropy and information. In *Proc. of the Fourth Berkeley Symposium on Math. Statistics and Probability*, volume 1, pages 547–561, 1961.

[RSSS17] Miruna Rosca, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Middle-product learning with errors. In *CRYPTO (3)*, volume 10403 of *Lecture Notes in Computer Science*, pages 283–297. Springer, 2017.

[RSW18] Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the ring-lwe and polynomial-lwe problems. In *EUROCRYPT (1)*, volume 10820 of *Lecture Notes in Computer Science*, pages 146–173. Springer, 2018.

[Rüc10] Markus Rückert. Lattice-based blind signatures. In *ASIACRYPT*, volume 6477 of *Lecture Notes in Computer Science*, pages 413–430. Springer, 2010.

[Saa18] Markku-Juhani O. Saarinen. Arithmetic coding and blinding countermeasures for lattice signatures - engineering a side-channel resistant post-quantum signature scheme with compact signatures. *J. Cryptogr. Eng.*, 8(1):71–84, 2018.

[Sch87] Claus-Peter Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.

[SE94] Claus-Peter Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Math. Program.*, 66:181–199, 1994.

[Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.

[SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, volume 5912 of *LNCS*, pages 617–635. Springer, 2009.

[vEH14] Tim van Erven and Peter Harremoes. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014.

[WW19] Y. Wang and M. Wang. Module-lwe versus ring-lwe, revisited. *IACR Cryptology ePrint Archive*, 2019:930, 2019.