



HAL
open science

Theoretical hardness of algebraically structured learning with errors

Katharina Boudgoust

► **To cite this version:**

Katharina Boudgoust. Theoretical hardness of algebraically structured learning with errors. Cryptography and Security [cs.CR]. Université Rennes 1, 2021. English. NNT : 2021REN1S064 . tel-03534254

HAL Id: tel-03534254

<https://theses.hal.science/tel-03534254v1>

Submitted on 19 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE DE DOCTORAT DE

L'UNIVERSITÉ DE RENNES 1

ÉCOLE DOCTORALE N° 601
*Mathématiques et Sciences et Technologies
de l'Information et de la Communication*
Spécialité : *Informatique*

Par

Katharina BOUDGOUST

Theoretical Hardness of Algebraically Structured Learning With Errors

Thèse présentée et soutenue à Rennes, le 16 novembre 2021
Unité de recherche : IRISA, UMR 6074

Rapporteurs avant soutenance :

Céline CHEVALIER Maîtresse de Conférences, Université Panthéon-Assas Paris 2
Frederik VERCAUTEREN Professeur Associé, Katholieke Universiteit Leuven

Composition du Jury :

Présidente :	Stéphanie DELAUNE	Directrice de Recherche, CNRS, Rennes, France
Examineurs :	Michel ABDALLA	Directeur de Recherche, CNRS, DIENS, Paris, France
	Shweta AGRAWAL	Professeure Assistante, IIT Madras, Inde
	Martin R. ALBRECHT	Professeur, Royal Holloway, University of London, Royaume-Uni
Rapporteurs :	Céline CHEVALIER	Maîtresse de Conférences, Université Panthéon-Assas Paris 2, France
	Frederik VERCAUTEREN	Professeur Associé, Katholieke Universiteit Leuven, Belgique
Dir. de thèse :	Adeline ROUX-LANGLOIS	Chargée de Recherche, Université de Rennes 1, CNRS, IRISA, France
Co-dir. de thèse :	Pierre-Alain FOUQUE	Professeur, Université de Rennes 1, France

Acknowledgments

As most projects in life, this thesis is not the fruit of work of one single person, but rather the output of a three years lasting collaboration with many competent and supportive people.

First of all, I would like to thank my two supervisors, Adeline Roux-Langlois and Pierre-Alain Fouque, who very generously accepted me as their PhD student at the research team EMSEC of the IRISA laboratory in Rennes. In particular, I very much appreciate the time and energy Adeline shared with me, helping me to navigate through the for me very new research field of lattice-based cryptography, as well as through French administration. I am very grateful that she gave me the opportunity to spend my PhD time in such comfortable research conditions with supportive people around me.

Secondly, I would like to thank Céline Chevalier and Frederik Vercauteren who kindly accepted to review my PhD manuscript. Alike, I very much appreciate that Michel Abdalla, Shweta Agrawal, Martin R. Albrecht and Stéphanie Delaune accepted to be, virtually or in person, part of the jury of my PhD defense.

Furthermore, I would like to articulate many thanks to my co-authors at the EMSEC team, Corentin Jeudy and Weiqiang Wen, who enriched my time in Rennes with innumerable discussions about challenging research questions. Moreover, a special thanks goes to Amin Sakzad and Ron Steinfeld who not only invited me to stay two months within their research group at the Monash University in Melbourne, but also continued our collaboration for many months afterwards.¹ Thanks also to my co-authors Shi Bai, Dipayan Das and Zhenfei Zhang with whom I collaborated for my very first research output.

Even though we don't share the same research field, I would like to thank my mentor Jean-Marie Bonnin, who accompanied me throughout the last two years with insightful discussions on the French research life. As said at the beginning, I am grateful for the good research environment I had at the IRISA research laboratory in Rennes. This does encompass the scientific aspects like regular seminars and workshops, the daily life at work among my research team EMSEC, but also more general aspects like the gender equality topics discussed and moderated by the Commission Égalité Femmes-Hommes.

I would like to finish my acknowledgments by expressing my enormous gratefulness towards the people, family and friends, who were present and close² to me in those last three years, especially during the very particular period of the Covid-19 pandemic. It is of immeasurable value to have kind people around me, supporting me in my way through Europe³ and helping me to navigate within this world.

¹Thanks to Damien Stehlé, who arranged a first get together with Amin and Ron in Lyon.

²Though, not necessarily geographically.

³The next stop will be at the Cryptography and Security group at Aarhus University in Denmark.

Contents

Notations and Abbreviations	vii
Résumé long en français	ix
Publications	xv
Introduction	xvii
1 Preliminaries	1
1.1 Algebraic Number Theory	2
1.1.1 Space H	3
1.1.2 Canonical Embedding, Trace and Norm	3
1.1.3 Coefficient Embedding and Rotation Matrix	4
1.1.4 Discrete Vandermonde Matrix	5
1.1.5 Modules, Ideals and Units	6
1.2 Lattice Theory	7
1.2.1 Euclidean Lattices	7
1.2.2 Structured Lattices	8
1.2.3 Computational Problems	9
1.3 Probability Theory	10
1.3.1 Gaussian Measures	10
1.3.2 Gaussians over Number Fields	13
1.3.3 Statistical Distance	16
1.3.4 Rényi Divergence	17
1.3.5 Leftover Hash Lemma	18
1.4 Learning With Errors	19
1.4.1 Learning With Errors (LWE)	19
1.4.2 Polynomial Learning With Errors (P-LWE)	20
1.4.3 Module Learning With Errors (M-LWE)	21
1.4.4 Middle-Product Learning With Errors (MP-LWE)	23
1.4.5 Learning With Rounding (LWR)	25
1.4.6 Practical Hardness of Learning With Errors	27
1.5 Cryptographic Notions	27
1.5.1 Public Key Encryption	27
1.5.2 Random Oracle Model	28

I	Theoretical Foundations	31
2	Binary Hardness of Module LWE	33
2.1	Introduction	33
2.1.1	Our Contributions	34
2.1.2	Related Work	36
2.1.3	Roadmap	38
2.2	Warm-up: A Simple Reduction	38
2.3	An Improved Reduction	41
2.3.1	First-is-errorless M-LWE	44
2.3.2	Extended M-LWE	45
2.3.3	Reduction to Binary M-LWE	50
2.4	Generalization to Larger Secrets	53
2.5	Choice of Embedding for Binary Secrets	55
2.5.1	Lagrange Basis	55
2.5.2	Lagrange Basis for Power-of-2 Cyclotomics	56
3	Classical Hardness of Module LWE	57
3.1	Introduction	57
3.1.1	Our Contributions	58
3.1.2	Related Work	61
3.1.3	Roadmap	62
3.2	Classical Reduction for Exponentially Large Moduli	62
3.2.1	Step 1: From Gaussian Decoding Problem to M-LWE	62
3.2.2	Step 2: Classical Hardness for Exponentially Large Modulus	67
3.3	Binary Hardness and Adapted Error Distribution	69
3.4	Modulus Reduction	70
3.4.1	The General Case	71
3.4.2	The Case of Power-of-2 Cyclotomics	72
4	Middle-Product Learning With Rounding	75
4.1	Introduction	75
4.1.1	Our Contributions	76
4.1.2	Related Work	77
4.1.3	Roadmap	77
4.2	Random Hankel Matrices	78
4.3	Middle-Product Learning With Rounding (MP-LWR)	81
4.4	Hardness of Computational MP-LWR	82
5	Partial Vandermonde Problems	89
5.1	Introduction	89
5.1.1	Our Contributions	90
5.2	Definition of Partial Vandermonde Problems	91
5.2.1	Partial Vandermonde Knapsack Problem (PV-Knap)	91
5.2.2	Partial Vandermonde Learning With Errors (PV-LWE)	93
5.2.3	PASS Problem	94
5.3	Equivalence of PV-Knap and PV-LWE	95
5.4	Efficient Computation of Partial Vandermonde Transform	97
5.4.1	Computation of the Partial Vandermonde Transformation	98
5.4.2	Computation of the Transposed Partial Vandermonde Transformation	98

II	Cryptographic Constructions	101
6	Encryption Based on Middle-Product LWR	103
6.1	Introduction	103
6.1.1	Our Contributions	104
6.1.2	Related Work	104
6.1.3	Roadmap	104
6.2	Reconciliation	104
6.3	Public Key Encryption Based on MP-LWR	106
6.3.1	Correctness	106
6.3.2	Security	107
6.4	Parameters and Security	111
6.4.1	Asymptotic Parameters	111
6.4.2	Concrete Security	112
7	PASS Encrypt	115
7.1	Introduction	115
7.1.1	Our Contributions	116
7.1.2	Related Work	116
7.1.3	Roadmap	117
7.2	PASS Encrypt	117
7.2.1	Correctness	118
7.2.2	Security	119
7.2.3	Homomorphic Properties	120
7.3	Security Analysis	121
7.3.1	Key Recovery Attack	122
7.3.2	Randomness Recovery Attack	123
7.3.3	Plaintext Recovery Using Hints Attack	123
7.3.4	Choice of Ring	124
7.3.5	Code-Based Attack	125
7.4	Concrete Parameters	125
7.4.1	Choice of the Number of Rows	126
7.4.2	Comparison	127
	Conclusion and Perspectives	129
	Bibliography	133

Notations and Abbreviations

We summarize here the notations used throughout the manuscript, listed in thematic order.

Symbol	Meaning
\mathbb{N}	Natural numbers (without 0)
\mathbb{C}, i	Field of complex numbers and the imaginary number i
$\Re(z), \bar{z}$	Real component and complex conjugate of a complex number $z \in \mathbb{C}$
\mathbb{R}, \mathbb{Z}	Field of real numbers and ring of rational numbers
\mathbb{Z}_q	Quotient ring of rational integers modulo some $q \in \mathbb{N}$
$[n]$	The set $\{1, \dots, n\}$ for some $n \in \mathbb{N}$
$\lfloor x \rfloor$	Closest integer to x
A	Bold capital letters denote matrices
$\mathbf{A}^T, \mathbf{A}^\dagger$	Transpose and Hermitian transpose of A
\mathbf{I}_n	The identity matrix of size $n \times n$
$[\mathbf{A} \mathbf{B}]$	Concatenation of two matrices A and B
$\mathbf{A} \otimes \mathbf{B}$	Kronecker product of two matrices A and B
a	Bold lowercase letters denote column vectors
$\mathbf{a} \circ \mathbf{b}$	Component-wise product of two vectors a and b
$\text{rev}(\mathbf{a})$	The vector $\mathbf{a} = (a_1, \dots, a_n)^T$ in reverse order, i.e., $\text{rev}(\mathbf{a}) = (a_n, \dots, a_1)^T$
$\text{diag}(\mathbf{a})$	The diagonal matrix whose diagonal entries are given by the vector a
$\{\mathbf{e}_j\}_{j \in [n]}$	The canonical basis of \mathbb{C}^n
$\mathbb{Z}[x]$	Ring of integer polynomials
$\mathbb{Z}[x]^{<n}$	Set of integer polynomials with degree less than n
$\ \mathbf{a}\ _\infty, \ \mathbf{a}\ _2$	Infinity and Euclidean norm of a vector a
$\ \mathbf{a}\ _{2,\infty}$	Hybrid norm of a vector over a number field K
$s_1(\mathbf{A})$	Largest singular value of a matrix A
$\ \mathbf{A}\ _2$	Spectral norm of a matrix A
$\ \mathbf{A}\ _{\max}$	Maximum norm of a matrix A
$\ \mathbf{A}\ _F$	Frobenius norm of a matrix A
$\text{GS}(\mathbf{A})$	Gram-Schmidt orthogonalization of A from left to right
K, R	A number field and its associated ring of integers
R^\times	Set of units (i.e., non-trivial invertible elements) of a ring R
\mathcal{M}, \mathcal{I}	Modules over a number field K and ideals over its ring of integers R
$\langle p \rangle = pR$	The principal ideal in a ring R generated by $p \in R$
R_p	The quotient ring $R/\langle p \rangle$
τ, σ, σ_H	The coefficient and the canonical embedding (H as a complex or real vector space)
$\bar{\sigma}$	The discrete canonical embedding over the quotient ring \mathbb{Z}_q
$\mathbf{V}, \bar{\mathbf{V}}$	The standard (over \mathbb{C}) and discrete (over \mathbb{Z}_q) Vandermonde matrix
$\bar{\mathbf{V}}_\Omega, \bar{\mathbf{V}}_{\Omega^c}$	The partial (discrete) Vandermonde matrix and its complement

$\delta_{j,k}$	Kronecker symbol, which equals 1 if $j = k$ and 0 otherwise
$\omega, O, \Theta, \Omega$	Standard Landau asymptotic notations
$U(S)$	Uniform distribution over a set S
$ S $	The cardinality of a set S
$x \leftarrow D$	Sampling an element $x \in S$ according to a distribution D over S
$\Delta(P; Q)$	Statistical distance between two probability distributions P and Q
$RD_2(P Q)$	Rényi divergence of order 2 between two probability distributions P and Q
D_s	Continuous Gaussian distribution of width s
$\mathcal{D}_{s,\Lambda}$	Discrete Gaussian distribution of width s over the lattice Λ
XOR	Bit-wise xor-operation of two binary vectors

Within this thesis we use the following abbreviations, listed in alphabetical order.

Abbreviation	Meaning
BDD	Bounded Distance Decoding
GDP	Gaussian Decoding Problem
HNF	Hermite Normal Form
IND-CPA	Indistinguishability Under Chosen Plaintext Attack
LHL	Leftover Hash Lemma
LWE	Learning With Errors
LWR	Learning With Rounding
M-LWE	Module Learning With Errors
MP-LWE	Middle-Product Learning With Errors
MP-CLWR	Middle-Product Computational Learning With Rounding
NTT	Number Theoretic Transform
OHCP	Oracle Hidden Center Problem
PASS	PASS Problem
PKE	Public Key Encryption
P-LWE	Polynomial Learning With Errors
PPT	Probabilistic Polynomial Time
PV-Knap	Partial Vandermonde Knapsack
PV-LWE	Partial Vandermonde Learning With Errors
PV-SIS	Partial Vandermonde Shortest Integer Solution
R-LWE	Ring Learning With Errors
ROM	Random Oracle Model
SIS	Short Integer Solution
SVP	Shortest Vector Problem

Résumé long en français

Historiquement, l'objectif de la cryptologie était de proposer des méthodes de chiffrement permettant de garantir la confidentialité des messages échangés entre deux personnes. Depuis le développement d'internet, d'autres défis sont également devenus importants, comme par exemple garantir l'authenticité d'une information en la signant numériquement, payer avec de l'argent électronique ou bien voter de manière anonyme lors d'élections électroniques. En utilisant un schéma cryptographique *prouvé sûr*, le fait qu'un message reste inintelligible, qu'une signature ne puisse pas être falsifiée ou qu'un vote reste anonyme est garanti par la difficulté d'un problème mathématique. Autrement dit, toute information divulguée pourrait être utilisée pour résoudre un problème mathématique réputé difficile. Cependant, la plupart des protocoles utilisés aujourd'hui sont menacés par l'arrivée d'ordinateurs quantiques de plus en plus puissants. Plus précisément, si nous avons accès à des ordinateurs qui utilisent la mécanique quantique au lieu de la mécanique classique, certains problèmes mathématiques deviendraient faciles à résoudre et, avec eux, de nombreux schémas cryptographiques actuels ne seraient plus sûrs [Sho97].

L'identification de protocoles cryptographiques supposés être protégés contre les algorithmes quantiques est un domaine de recherche actif appelé la *cryptographie post-quantique*. Dans ce contexte, l'institut américain National Institute of Standards and Technology (NIST) a lancé en 2016 un processus de normalisation post-quantique [NIS] pour les schémas de chiffrement à clé publique et les signatures numériques. Plusieurs pistes qui semblent résister aux attaques quantiques ont été proposées. Pour n'en citer que quelques-unes, il existe des hypothèses à base de réseaux euclidiens, des codes aléatoires, des fonctions de hachage ou des systèmes multivariés. Les schémas qui reposent sur des problèmes sur les réseaux euclidiens constituent une catégorie très prometteuse. Ce domaine de recherche est appelé *cryptographie à base de réseaux* et offre de multiples avantages, tels que l'efficacité, la versatilité et des garanties de sécurité élevées. Pour illustrer son rôle prépondérant, notons que dans le processus de normalisation en cours géré par le NIST, parmi les candidats restants au dernier tour, 3 des 4 mécanismes d'encapsulation de clés (permettant le chiffrement) et 2 signatures numériques sur 3 sont basés sur des réseaux. Toutes les contributions présentées dans ce manuscrit s'inscrivent dans le domaine de la cryptographie à base de réseaux.

Au cœur de la plupart des schémas reposant sur les réseaux euclidiens, et en particulier au centre de cette thèse, se trouve un problème de calcul, celui de l'apprentissage avec erreurs, souvent abrégé LWE (acronyme de l'anglais Learning With Errors). De manière informelle, le problème LWE demande de résoudre un système d'équations linéaires bruitées sur les entiers rationnels. Plus formellement, pour des entiers positifs n, m et q , soit $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ une matrice tirée selon la distribution uniforme sur \mathbb{Z}_q . De plus, on considère deux distributions de probabilités. La première, notée ψ_e , est sur \mathbb{Z} et utilisée pour générer un vecteur $\mathbf{e} \in \mathbb{Z}^m$ d'une norme euclidienne petite avec forte probabilité. Souvent, on choisit une distribution gaussienne pour ψ_e . La deuxième distribution de probabilités, désignée par ψ_s , est sur \mathbb{Z}_q^n et définit un vecteur $\mathbf{s} \in \mathbb{Z}_q^n$. Nous appelons ces deux vecteurs le *bruit/erreur* \mathbf{e} et le *secret* \mathbf{s} . Une instance

de LWE est donnée par le couple $(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e} \bmod q)$, tandis que le bruit et le secret restent cachés. La variante calculatoire de LWE demande de trouver le secret \mathbf{s} et la version décisionnelle demande de distinguer cette instance d’une instance (\mathbf{A}, \mathbf{b}) où le vecteur \mathbf{b} a été tiré de la distribution uniforme sur \mathbb{Z}_q^m . Notons que sans le bruit \mathbf{e} , il s’agirait d’un système standard d’équations linéaires qui peut être résolu en temps polynomial par élimination gaussienne. Seul le bruit supplémentaire rend la résolution du système d’équations difficile. Pour un certain facteur d’approximation qui dépend des paramètres ψ_s, ψ_e, m, n et q , on peut prouver que LWE est au moins aussi difficile que de trouver approximativement le vecteur le plus court dans un réseau euclidien [Reg05, Pei09, BLP⁺13, PRS17a]. Ceci est un problème supposé difficile pour lequel aucun algorithme efficace n’est connu lorsque la dimension du réseau est suffisamment grande.

Malheureusement, les protocoles cryptographiques qui s’appuient sur la difficulté de LWE sont intrinsèquement inefficaces en raison de la taille des clés publiques qui contiennent généralement la matrice $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, où n et m sont au moins aussi grands que le paramètre de sécurité. En outre, les opérations de base sont des produits matrice-vecteur sur \mathbb{Z}_q qui nécessitent un temps de calcul quadratique en la dimension n . Pour améliorer l’efficacité des protocoles, des variantes structurées de LWE ont été proposées, comme par exemple [SSTX09, LPR10, LS15].

L’une d’entre elles est le problème Module Learning With Errors (M-LWE), introduit par Brakerski et al. [BGV12] et étudié en détails par Langlois et Stehlé [LS15]. Au lieu de considérer des équations linéaires bruitées sur les entiers \mathbb{Z} , les équations linéaires sont maintenant définies algébriquement sur un anneau de polynômes à coefficients dans \mathbb{Z} . Intuitivement, l’amélioration de l’efficacité vient du fait que chaque polynôme de cet anneau définit une matrice structurée de multiplication. Ainsi, nous pouvons remplacer la matrice complètement aléatoire de LWE par une certaine matrice structurée pour M-LWE. De plus, dans ce cas le produit matrice-vecteur peut être calculé en temps quasi-linéaire. Depuis son introduction en 2012, M-LWE jouit d’une popularité croissante car il offre un compromis fin entre sécurité concrète et efficacité. Dans le cadre du processus de normalisation du NIST [NIS], plusieurs candidats au troisième tour s’appuient sur la difficulté de M-LWE, par exemple, le schéma de signature Dilithium [DKL⁺18] et le mécanisme d’encapsulation des clés Kyber [BDK⁺18].

Contributions

La recherche présentée dans ce manuscrit se concentre sur les variantes structurées du problème LWE et leur utilisation en cryptographie post-quantique. Dans ce qui suit, nous décrivons les contributions présentées qui peuvent être divisées en deux catégories différentes. La première partie contient des contributions liées à une meilleure compréhension des fondations théoriques de la cryptographie à base de réseaux et la deuxième partie comprend la conception de deux schémas de chiffrement efficaces, dont la sécurité est basée sur des problèmes de réseaux structurés.

Fondations théoriques

Comme nous l’avons expliqué précédemment, les schémas cryptographiques efficaces sont basés sur des variantes structurées de LWE, comme par exemple M-LWE. À de nombreux égards, le problème plus récent M-LWE ressemble à son homologue non structuré et plus étudié LWE. Néanmoins, certaines propriétés importantes qui ont été prouvées pour LWE, n’ont pas (encore) été démontrées pour M-LWE. Nous pensons qu’une étude rigoureuse du problème M-LWE est essentielle pour approfondir notre confiance dans les hypothèses de difficulté qui sont faites dans les schémas cryptographiques actuellement proposés pour la normalisation. Dans ce qui suit, nous présentons deux résultats qui répondent à deux de ces insuffisances concernant le problème M-LWE.

Secret binaire

Comme première contribution, nous prouvons dans le Chapitre 2 que M-LWE reste difficile à résoudre, même si la distribution du secret ψ_s donnant le vecteur s sur l’anneau est modifiée. Au lieu d’utiliser la distribution uniforme qui conduit à des polynômes à grands coefficients (c’est-à-dire dans \mathbb{Z}_q), nous pouvons également utiliser une distribution de probabilités qui donne des polynômes à coefficients binaires (c’est-à-dire dans $\{0, 1\}$). Nous désignons ce dernier problème par **bin-M-LWE**. Il est particulièrement intéressant car il augmente l’efficacité des calculs, permet des techniques efficaces de changement du rang et du modulo [BLP⁺13, AD17a, WW19] et est nécessaire dans les schémas de chiffrement totalement homomorphes comme dans [DM15, CGGI16]. Pour le cas non structuré, des réductions de LWE à **bin-LWE** avec une dimension et une borne d’erreur légèrement augmentées ont été montrées [GKPV10, BLP⁺13, Mic18]. Cependant, tous les résultats existants sont sur les entiers \mathbb{Z} et non sur l’anneau des entiers d’un corps de nombres. Par conséquent, comme l’indique Micciancio dans la conclusion de [Mic18], une question ouverte importante est de savoir si des résultats similaires s’appliquent aux variantes structurées, en particulier à M-LWE. Dans cette thèse, nous progressons vers la résolution de ce problème en prouvant la difficulté de **bin-M-LWE**, si le rang du module est (super-)logarithmique dans le degré du corps de nombres sous-jacent. Plus précisément, nous présentons deux preuves différentes pour obtenir ce résultat, qui diffèrent légèrement dans leurs conditions, les paramètres obtenus et les variantes concrètes du problème. De plus, nous montrons que les deux réductions peuvent être généralisées à une distribution du secret plus large.

Réduction classique

Comme deuxième contribution portant sur la difficulté théorique de M-LWE, nous renforçons dans le Chapitre 3 son lien étroit avec les problèmes sur les réseaux euclidiens structurés en prouvant une réduction classique. Auparavant, seule une réduction quantique était connue [LS15]. Une telle réduction nécessite des ordinateurs quantiques, qui sont extrêmement puissants, coûteux à construire et pas encore disponibles à grande échelle. Pour le cas non structuré, une preuve de difficulté classique de LWE est donnée par Brakerski et al. [BLP⁺13]. Sur le plan général, nous suivons la même structure que dans [BLP⁺13] en adaptant de façon rigoureuse tous les résultats nécessaires au contexte des modules. Plus en détail, nous avons besoin de trois ingrédients. Premièrement, il nous faut une réduction classique d’un problème sur les réseaux modules au problème M-LWE avec un modulo de taille exponentielle, que nous obtenons en adaptant le pendant de LWE de Peikert [Pei09] que nous combinons avec un résultat plus récent de Peikert et al. [PRS17a]. Comme deuxième composant, nous avons besoin de la difficulté de M-LWE en utilisant un secret binaire, que nous prouvons dans le Chapitre 2. Enfin, une technique de réduction du paramètre de modulo est nécessaire, où nous pouvons utiliser la technique de changement du rang et du modulo montrée par Albrecht et Deo [AD17a]. En assemblant soigneusement ces trois ingrédients, nous obtenons finalement la réduction classique.

Middle-Product Learning With Rounding

Après avoir étudié le problème M-LWE qui sert d’hypothèse de difficulté sous-jacente aux schémas de chiffrement et aux signatures numériques pratiques, nous nous intéressons dans le Chapitre 4 à une toute nouvelle version structurée de LWE. La motivation principale de ce chapitre est de combiner deux variantes existantes de LWE, la variante Middle-Product Learning With Errors (MP-LWE) [RSSH17] et la variante Learning With Rounding (LWR) [BPR12], afin de définir un nouveau problème qui bénéficie de leurs avantages respectifs. À cette fin, nous introduisons une nouvelle hypothèse de difficulté que nous appelons le problème Middle-Product Compu-

tational Learning With Rounding (MP-CLWR). D'une part, MP-CLWR utilise les arrondis de manière similaire à LWR et évite ainsi l'échantillonnage de l'erreur gaussienne. D'autre part, la difficulté de MP-CLWR ne dépend pas d'un anneau spécifique, mais est garantie par la difficulté d'une variante structurée de LWE (appelée P-LWE) pour un ensemble d'anneaux exponentiellement grand. Ainsi, l'hypothèse MP-CLWR jouit des propriétés souhaitées, à la fois de l'avantage de sécurité de MP-LWE et de l'avantage de simplicité de LWR. Concernant la difficulté de ce nouveau problème, nous prouvons que MP-CLWR est au moins aussi difficile que MP-LWE, qui est lui-même basé sur la difficulté des problèmes sur les réseaux euclidiens structurés.

Problèmes liés à la matrice de Vandermonde partielle

Dans les chapitres précédents, nous avons étudié des variantes de LWE dont la difficulté est garantie par l'impossibilité présumée de résoudre des problèmes dans les pires cas sur des réseaux euclidiens structurés. Dans le Chapitre 5, nous étudions une famille différente de problèmes sur des réseaux structurés qui est liée à la transformée de Vandermonde discrète et qui permet des protocoles simples et efficaces au détriment de l'absence de connexion connue entre les pires cas et les cas moyens. L'hypothèse constituant la base de ce chapitre est qu'il n'existe ni d'algorithme classique ni quantique en temps polynomial qui récupère un polynôme de petits coefficients en ayant seulement accès à une liste partielle de sa transformée de Vandermonde discrète. Bien sûr, il est essentiel de préciser à quel point les coefficients du polynôme doivent être petits et de quelle taille est la liste partielle que nous fournissons. Par analogie avec le problème du sac à dos standard, nous appelons ce problème le problème de Partial Vandermonde Knapsack (PV-Knap). Sa version homogène est appelée Partial Vandermonde SIS (PV-SIS), et peut être vue comme une variante spécifique du problème standard Short Integer Solution.

En 2015, Hoffstein et Silverman [HS15] présentent PASS Encrypt, un schéma de chiffrement à clé publique dont les briques de construction sont liées au problème PV-Knap. Leur schéma de chiffrement est très efficace et remplit des propriétés homomorphiques additives et multiplicatives, ce qui en fait un point de départ naturel pour la conception de primitives cryptographiques efficaces. Malheureusement, l'un des principaux inconvénients de PASS Encrypt est qu'aucune preuve de sécurité n'a été donnée dans [HS15]. Dans ce chapitre, nous progressons dans la compréhension des hypothèses de difficulté nécessaires pour prouver la sécurité de PASS Encrypt. D'abord, nous élargissons le paysage des problèmes qui utilisent la transformation partielle de Vandermonde en définissant une nouvelle variante de LWE, appelée Partial Vandermonde Learning With Errors (PV-LWE). Par la suite, nous montrons l'équivalence de PV-Knap et de PV-LWE en exploitant la même connexion de dualité que nous connaissons pour les problèmes Knapsack et LWE standards. Comme notre principale motivation est de fournir une preuve de sécurité pour PASS Encrypt, nous définissons une variante du PV-Knap, que nous appelons le problème PASS. Ce problème sert (avec la version décisionnelle de PV-Knap) d'hypothèse de difficulté sous-jacente pour (une version légèrement modifiée de) PASS Encrypt. Nous présentons le schéma ainsi que la preuve de sécurité plus tard dans le Chapitre 7.

Constructions cryptographiques

Dans la première partie du manuscrit, nous introduisons deux nouvelles hypothèses de difficulté liées aux problèmes sur les réseaux euclidiens structurés. Afin de montrer comment se servir de ces hypothèses pour la cryptographie, nous construisons deux schémas de chiffrement à clé publique dans la deuxième partie de ce travail. Nous prouvons que les deux schémas sont corrects et sûrs en supposant des hypothèses de calcul explicitement énoncées.

Chiffrement basé sur Middle-Product LWE

Pour montrer l'utilité cryptographique de notre nouvelle hypothèse de difficulté MP-CLWR du Chapitre 4, nous présentons au Chapitre 6 un schéma de chiffrement dont la sécurité est basée sur MP-CLWR. Sa conception s'inspire simultanément de deux schémas existants : d'une part, du schéma de Roşca et al. [RSSS17] dont la sécurité est basée sur MP-LWE, et d'autre part, du chiffrement de Chen et al. [CZZ18] dont la sécurité est basée sur la variante structurée de LWR. Comme pour [CZZ18], nous utilisons un mécanisme dit de réconciliation, qui permet de déchiffrer correctement. Nous prouvons la sécurité de notre schéma en nous basant sur la difficulté de MP-CLWR. Comme c'est le cas pour les schémas qui utilisent des arrondis, pendant la génération de la clé publique, nous n'avons besoin que d'arrondir le produit médian de deux polynômes au lieu d'échantillonner l'erreur gaussienne, ce qui le rend plus facile à mettre en œuvre que le schéma de [RSSS17]. Simultanément, tout en garantissant un niveau de sécurité égal, nous obtenons les mêmes tailles asymptotiques de clé et de chiffrements que le schéma de [RSSS17]. Enfin, nous analysons la sécurité concrète de notre schéma en examinant les meilleures attaques actuellement connues contre lui.

PASS Encrypt

Comme contribution finale, nous présentons dans le Chapitre 7 une version modifiée du schéma de chiffrement PASS Encrypt [HS15] ainsi qu'une preuve de sécurité basée sur le problème PV-Knap et une variante particulière de celui-ci, que nous appelons le problème PASS, tous les deux étudiés dans le Chapitre 5. Ce dernier problème capture le fait qu'un message chiffré de PASS Encrypt consiste en plusieurs transformées de Vandermonde partielles des éléments *liés*. Par conséquent, un*e attaquant*e qui réussit contre PV-Knap peut être utilisé*e pour gagner le jeu de sécurité de PASS Encrypt, mais le contraire n'est pas vrai. Ce problème n'a pas été abordé auparavant dans la version originale de PASS Encrypt par Hoffstein et Silverman [HS15].

Nous modifions légèrement le schéma et donnons une preuve que le déchiffrement est correct pour des paramètres bien choisis et une preuve de sécurité, en supposant la difficulté de PV-Knap et de PASS. Nous fournissons également une analyse fine de la sécurité du schéma en montrant une nouvelle attaque que nous appelons Plaintext Recovering Using Hints attack, qui prend en compte la structure de PASS Encrypt. À cette fin, nous utilisons un travail récent de Dachman-Soled et al. [DDGR20] pour analyser les instances de LWE, où des indices supplémentaires sur le secret et/ou l'erreur sont donnés. Nos estimations de complexité pour cette attaque montrent qu'elle ne réduit pas la complexité en dessous de celle des attaques précédemment connues sur PASS, ce qui augmente notre confiance dans la sécurité revendiquée de PASS Encrypt. Nous concluons le chapitre en fournissant des exemples concrets de paramètres et nous comparons notre schéma avec deux autres schémas efficaces dont la sécurité est basée sur des problèmes de réseaux structurés.

Publications

In the following, we list the peer-reviewed publications in the proceedings of international conferences that build the basis of this thesis. In cryptology, publications are mainly done in conference proceedings, and few in journals, and the authors are listed in alphabetical order.

[BJRW21] On the Hardness of Module-LWE with Binary Secret [\[HAL\]](#) Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois and Weiqiang Wen. Published in the proceedings of [CT-RSA 2021](#) (pp. 503-526).

[BJRW20] Towards Classical Hardness of Module-LWE: The Linear Rank Case [\[HAL\]](#) Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois et Weiqiang Wen. Published in the proceedings of [Asiacrypt 2020](#) (pp. 289-317).

[BBD⁺19] Middle-Product Learning with Rounding Problem and its Applications [\[HAL\]](#) Shi Bai, Katharina Boudgoust, Dipayan Das, Adeline Roux-Langlois, Weiqiang Wen and Zhenfei Zhang. Published in the proceedings of [Asiacrypt 2019](#) (pp. 55-81).

Introduction

The word *cryptology* is composed of two ancient Greek words, *κρυπτος* (*kryptos*, secret) and *γραφειν* (*graphein*, to write), capturing the association that most people have when they hear about cryptography: exchanging messages in such a way that no one other than the recipient can read them. However, this only covers one of many different aspects of cryptography. From a general perspective, we can state three basic goals that cryptography aims to attend: confidentiality, authenticity and integrity. Whereas the first one corresponds to keeping the message secret, the second one guarantees a recipient that the message was indeed sent by the claimed sender, and the last one makes sure that the content of the message has not been altered. One can well imagine that it makes a difference to obtain a check of 1.000 Euro or 100 Euro, whereas the two numbers are only one 0 far away from each other. Modern cryptography does not only make it possible to exchange digital messages in a secure way, as for instance when sending encrypted emails via OpenPGP or using a messenger application with end-to-end encryption such as Signal. It also provides solutions to generally authenticate content on the internet, to establish secure connections between our home computer and the world wide web (e.g., the TLS protocol), to manage online banking, to do online shopping, and even gives us the possibility to securely vote online. This said, modern cryptography has become indispensable and omnipresent in the everyday life of a citizen in the 21st century.

But what exactly do we mean by *modern* cryptography? For a very long time, cryptography was done in a secret way, without publicly telling the rules. It was mainly used by governments and the military, who had the political and financial power to actually deploy cryptography. A now very popular example is the encryption machine *Enigma* used by the German military and marines during World War 2, whose breaking by the Allies is estimated to have shortened the war by several years and was kept secret a long time afterwards.⁴ The area of modern cryptography was marked by the publications of Diffie and Hellman [DH76] and of Merkle [Mer78], introducing the concept of *public key cryptography*.⁵ Its underlying principle is to equip everyone with a pair of cryptographic keys. The first one, referred to as the secret key, should be kept hidden. The second one is called the public key and can be made generally available. In the context of encryption, the public key allows anyone to encrypt a message, whereas only the owner of the corresponding secret key can decrypt the encrypted message. The main advantage of such a public key encryption is that the two persons who want to safely communicate with each other don't have to (physically) meet beforehand. As opposed to public key cryptography, there also exists a *secret key* counterpart, where both participants have to share a common secret

⁴As a first and accessible approach to the topic, we recommend the film *The Imitation Game* by Morten Tyldum (2014). Also very interesting is the role many women, called *code girls*, played during World War 2 (but already before) to decrypt messages sent by the enemy, here the book *Code Girls* by Liza Mundy (2017) gives many interesting insights for the US and the UK.

⁵Roughly at the same time, employees of the UK Government Communications Headquarters (GCHQ) made similar discoveries which were, however, only declassified in 1997 and thus not accessible for public research at that time.

key in order to execute the cryptographic scheme. A very popular example of a secret key encryption scheme is the AES block cipher designed by Daemen and Rijmen [DR00]. Such schemes are usually more efficient than their public key analogues, but require the exchange of secret credentials to initiate them. In practice, we generally use a public key protocol to safely exchange key material to enable an efficient secret key protocol afterwards.

Preferably, the public key encryption scheme comes with a *security proof*, guaranteeing that a message remains completely unintelligible as long as an explicitly formulated assumption is true. Typically, such an assumption states that some mathematical problem is difficult to solve. In other words, any disclosed information on the encrypted message could be used to solve a mathematical problem deemed unsolvable in a reasonable amount of time. One example of such a mathematical problem is the task of factoring a large number. A common conjecture is that there exists no polynomial-time algorithm that factors a number which is the product of two equally large primes. Another mathematical problem used in cryptography is the discrete logarithm problem, which asks to find the discrete logarithm with respect to a given basis in a finite group. Most of the current cryptographic schemes used in practice rely on variants of the factoring or the discrete logarithm problem, as for instance the Diffie-Hellman key exchange protocol [DH76], the RSA digital signature and encryption scheme [RSA78] or the ElGamal encryption [Gam84].

Unfortunately, those cryptographic schemes are threatened as the underlying assumptions may become vacuous once we possess *quantum* computers which, as opposed to classical computers, operate with quantum mechanics and can execute quantum algorithms. More precisely, in 1994 Shor [Sho97] designed a quantum algorithm that solves the discrete logarithm problem and the factorization problem in polynomial-time.⁶ It is important to emphasize that quantum computers do not yet exist on a large scale. In 2020, Google's quantum computer Hummingbird was able to handle 65 quantum bits (qubits) and the company hopes to build a quantum computer which manages roughly 1000 qubits in 2023 [Gam]. Note that different works hint towards the need of at least 1 million qubits (in run-time or memory) to factor a 2048 bits RSA integer [GE21, GS21]. However, it is crucial to start today to think about alternatives for a future, where quantum computers will be widely available, as safely replacing current cryptographic protocols will take years. Furthermore, a malicious person could start collecting encrypted messages today in order to decrypt them later, once they⁷ possess powerful quantum computers. Identifying and constructing cryptographic schemes that are assumed to be secure against quantum algorithms is an active research area called *post-quantum cryptography*. Within this context, the US National Institute of Standards and Technology (NIST) has initiated in 2016 a standardization process [NIS] for public-key encryption schemes and digital signatures which can be used to replace the current standards once large-scale quantum computers are available. There have been several directions proposed to build public key cryptography based on different mathematical hard problems, that seem to withstand quantum attacks. To name a few, there are assumptions based on lattices, random codes, hash functions, isogenies between elliptic curves or multivariate systems.

One very promising category of quantum-resistant candidates are schemes that rely on the hardness of problems on Euclidean lattices. This research area is called *lattice-based cryptography* and emerged at the end of the 1990s through the works of two different research lines. On the one hand, there have been proposals benefiting from strong theoretical connections to presumed hard worst-case lattice problems by Ajtai and Dwork [Ajt96, AD97] and Regev [Reg05]. On the other hand, very efficient schemes basing their security on average-case structured lattice

⁶The impact of quantum computers on secret key cryptography, for instance AES, is so far limited to quadratic speed-ups using Grover's algorithm [Gro97]. This effect can easily be mitigated by doubling the key sizes.

⁷Throughout the thesis, the neutral singular pronouns *they/their* are used in order to keep the language as inclusive as possible. See also <https://www.acm.org/diversity-inclusion/words-matter>




Figure 1: A lattice Λ in the 2-dimensional space \mathbb{R}^2 with basis $\mathbf{B} = (\mathbf{b}_1, \mathbf{b}_2)$ and $\widehat{\mathbf{B}} = (\widehat{\mathbf{b}}_1, \widehat{\mathbf{b}}_2)$, where the volume of the area framed by both bases (colored in light orange and light gray) is the same. The minimum of the lattice is given by $\lambda_1(\Lambda) = \|\mathbf{b}_1\|_2 = \|\mathbf{b}_2\|_2$.

problems have been introduced, the most popular among them is the NTRU encryption scheme by Hoffstein et al. [HPS98]. Since then, the field attracted more and more research interest, now estimated as one of the most mature candidates to replace standard cryptographic systems in a quantum world. To illustrate its leading role, note that in the on-going standardization process run by the NIST, among the remaining candidates in the final round, 3 out of 4 key encapsulation mechanisms (enabling encryption) and 2 out of 3 digital signatures are based on lattices.⁸ All contributions presented in this thesis are placed within the field of lattice-based cryptography.

Lattices are mathematical objects that play an important role in many different areas such as number theory, geometry and group theory. Informally speaking, a lattice is composed of points that are arranged in a periodic structure in the n -dimensional space \mathbb{R}^n , as illustrated in Figure 1. More formally, a set of n linear independent vectors $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n) \in \mathbb{R}^{n \times n}$ (called a basis) defines a lattice Λ given by

$$\Lambda = \left\{ \sum_{j=1}^n x_j \mathbf{b}_j : x_j \in \mathbb{Z} \right\}.$$

A lattice can have many different bases, but they all frame an area of the same volume, called the fundamental domain. Another invariant of the lattice Λ , which doesn't depend on the chosen basis to represent it, is the Euclidean norm of any smallest non-zero vector in Λ , called its first minimum and denoted by $\lambda_1(\Lambda)$.

One of the most studied problems related to Euclidean lattices is the Shortest Vector Problem (SVP), which asks for a given lattice Λ to find a shortest non-zero vector in Λ , i.e., a vector of norm $\lambda_1(\Lambda)$. A relaxed version of this problem is to find a shortest non-zero vector up to some approximation factor $\gamma \geq 1$, i.e., a vector of norm at most $\gamma \cdot \lambda_1(\Lambda)$, denoted by SVP_γ . It also possesses a promise decision variant, where we only ask, for a given parameter δ , to decide whether a given lattice is a YES instance, where $\lambda_1(\Lambda) \leq \delta$, or a NO instance, where $\lambda_1(\Lambda) > \gamma \cdot \delta$. We denote this problem by GapSVP_γ and illustrate it in Figure 2 for the two-dimensional case. Another lattice problem of interest is the approximate Shortest Independent Vector Problem (SIVP_γ).

⁸<https://csrc.nist.gov/Projects/post-quantum-cryptography/round-3-submissions>

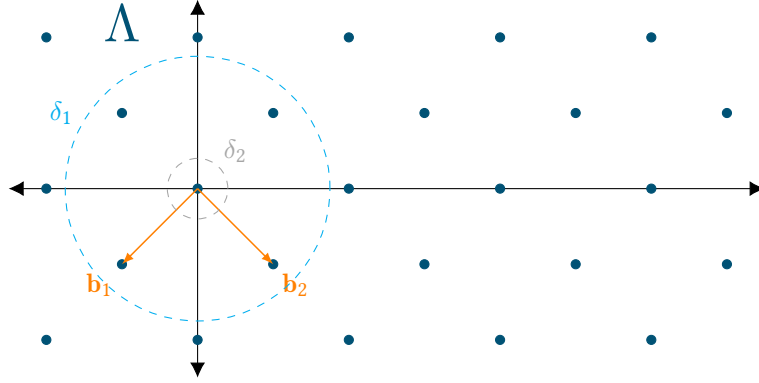


Figure 2: Given a lattice Λ with basis $\mathbf{B} = (\mathbf{b}_1, \mathbf{b}_2)$ and two positive reals δ_1 and δ_2 . It yields $\lambda_1(\Lambda) = \|\mathbf{b}_1\|_2$, $\lambda_1(\Lambda) \leq \delta_1$ and $\lambda_1(\Lambda) > \gamma \cdot \delta_2$ for $\gamma = 2$. In other words, Λ defines a YES instance for δ_1 and a NO instance for δ_2 for the problem GapSVP_2 .

Given an n -dimensional lattice Λ , SIVP_γ asks to find n linearly independent vectors of Λ that are shorter than $\gamma \cdot \lambda_n(\Lambda)$. Here, $\lambda_n(\Lambda)$ denotes the n -th minimum of Λ , i.e., the smallest number such that there exist n linearly independent vectors in Λ of norm less than it. In 1982, Lenstra, Lenstra and Lovász [LLL82] designed the now very popular LLL algorithm, that solves in polynomial time SVP_γ for γ exponentially large in the lattice dimension. Later in 1987, Schnorr showed a trade-off between running time and approximation factor which can be achieved by an algorithm solving SVP_γ [Sch87]. In practice, it is implemented by the BKZ algorithm by Schnorr and Euchner [SE94], which can be seen as a heuristic variant of Schnorr's algorithm. Following this trade-off, the best known algorithm to solve SVP_γ with γ polynomial in the lattice dimension n has an exponential running time of $2^{\tilde{O}(n)}$ and, conversely, the best known algorithm to solve SVP_γ with polynomial running-time can only achieve an exponential approximation factor γ of $2^{\tilde{O}(n)}$. Here, the term $\tilde{O}(n)$ designs the big O notation which hides logarithmic factors in n . Compared to the other two lattice problems, SVP_γ is no easier than SIVP_γ [SD16] which is itself no easier than GapSVP_γ [Ban93].⁹ The above leads to the following conjecture which forms the starting point of lattice-based cryptography. Note that all asymptotic statements are with respect to the lattice dimension n , unless we state it otherwise.

Conjecture 1: There is no polynomial-time classical or quantum algorithm that approximates the lattice problems SVP_γ , GapSVP_γ or SIVP_γ to within polynomial factors γ .

Despite their assumed quantum-resistance, those problems seem unlikely to directly serve for the construction of cryptographic primitives. This is because their definition relies on arbitrary lattices, what we commonly call *worst-case* problems, as they are in general not hard to solve for any lattice, but assumed to be hard to solve in the worst-case. When designing cryptographic schemes, however, we usually need the hardness of random instances of some problem, what we call *average-case* problems. This challenge was solved with the help of intermediate lattice prob-

⁹Note that the approximation factors are not preserved through the reductions. In the reduction from GapSVP_γ to SIVP_γ it is mapped from γ to $n\gamma$, where n is the lattice dimension. And in the reduction from SIVP_γ to SVP_γ it is mapped from γ to $\sqrt{n}\gamma$.

lems, namely the Short Integer Solution (SIS) [Ajt96] and Learning With Errors (LWE) [Reg05] problems, which are formulated as average-case problems, making them suitable for cryptography. Astonishingly, these intermediate lattice problems have been shown to be at least as hard to solve as some worst-case lattice problems, such as GapSVP_γ or SIVP_γ , for suitable parameter choices. Thus, Conjecture 1 which states that SIVP_γ and GapSVP_γ are classically and quantumly intractable implies that SIS and LWE are also classically and quantumly intractable.

In the following, we focus our discussion on the Learning With Errors (LWE) problem, which was introduced by Regev [Reg05] in his pioneering work. It is not only at the heart of most lattice-based schemes, but also builds the core element of this thesis. Informally speaking, an instance of LWE is a system of noisy linear modular equations over the rational integers \mathbb{Z} and it can be formulated either as a search or as a decision problem. Its search variant asks to find a solution to this system, whereas its decision version asks to distinguish such noisy linear modular equations from uniformly random ones. More formally, for positive integers n, m and q , let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ be a matrix sampled uniformly at random over the quotient ring \mathbb{Z}_q . Further, we consider two probability distributions. The first one, denoted by ψ_e , is over \mathbb{Z} and provides a vector $\mathbf{e} \in \mathbb{Z}^m$ whose Euclidean norm is small with respect to q with high probability. A common choice for ψ_e is a discrete Gaussian distribution. The second one, denoted by ψ_s , is over \mathbb{Z}_q^n and serves for sampling a vector $\mathbf{s} \in \mathbb{Z}_q^n$. In the following, you can think of ψ_s as the uniform distribution. We commonly refer to \mathbf{s} as the *secret* and to \mathbf{e} as the *error* or *noise*. An instance of LWE is given by $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$, where both \mathbf{s} and \mathbf{e} should be kept hidden. The search problem requires to find \mathbf{s} (or equivalently \mathbf{e}) as illustrated in Figure 3 and the decision problem asks to distinguish such an instance from (\mathbf{A}, \mathbf{b}) , where \mathbf{b} was sampled uniformly at random over \mathbb{Z}_q^m . Note that without the error term \mathbf{e} this would be a standard system of modular linear equations which can be solved in polynomial time by Gaussian elimination. Only the additional noise makes solving the system of equations very difficult. For some approximation factor γ depending on the parameters ψ_s, ψ_e, m, n and q , LWE is shown to be no easier to solve than worst-case lattice problems such as GapSVP_γ or SIVP_γ [Reg05, Pei09, BLP⁺13, PRS17a].

The worst-to-average case reduction from SIVP_γ to LWE [Reg05] only works when the error \mathbf{e} is sampled from a Gaussian distribution. This requirement is closely related to the nature of the reduction which reduces standard lattice problems such as SIVP_γ to the problem of sampling elements from a narrow discrete Gaussian distribution. However, Gaussian sampling procedures are in general costly, difficult to implement and vulnerable to side-channel attacks, e.g., [DB15, BHL16, Pes16, Saa18]. In 2012, Banerjee et al. [BPR12] introduce a deterministic variant of LWE, namely the Learning With Rounding (LWR) problem. Instead of adding a random error vector \mathbf{e} to the matrix-vector product $\mathbf{A}\mathbf{s}$ over \mathbb{Z}_q , we deterministically round the product modulo some smaller integer p . Informally, the noise of LWE helps to *hide* the low-order bits of $\mathbf{A}\mathbf{s}$, whereas the rounding simply *cuts* them off. It is shown, that LWR is no easier than LWE for a specific parameter setting [BPR12, AKPW13, BGM⁺16, AA16].

The decision variant of LWE serves as the hardness assumption for a wide range of provably secure cryptographic primitives, from the basic ones, such as Public Key Encryption (e.g., [Reg05, MP12]), to the most advanced ones, such as Fully Homomorphic Encryption (e.g., [BGV12, BV14, DM15]) or Non-Interactive Zero-Knowledge proof systems (e.g., [PS19]).

Unfortunately, the cryptographic protocols relying on the hardness of LWE are inherently inefficient due to the size of the public keys which usually contain the matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, where n is at least as large as the security parameter and m is the number of samples which is usually larger than $n \log(n)$. Furthermore, the basic operations are matrix-vector products over \mathbb{Z}_q which require a computation time quadratic in the dimension n . To improve the efficiency, structured variants of LWE have been proposed, e.g., [SSTX09, LPR10, LS15].

One of them is the Module Learning With Errors (M-LWE) problem, introduced by Brakerski

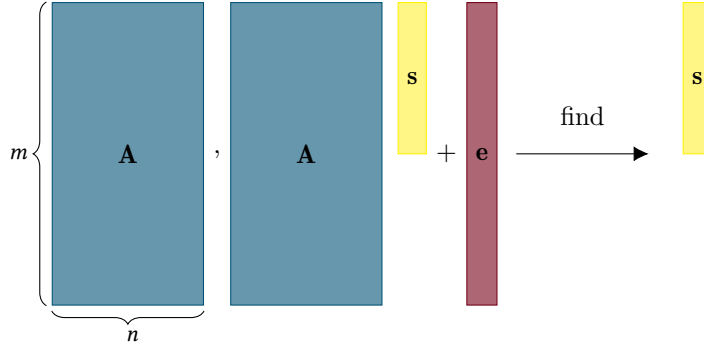


Figure 3: The LWE problem in its search variant. The number of rows m of \mathbf{A} can be seen as the number of LWE samples and the number of columns n of \mathbf{A} defines the dimension of the LWE problem.

et al. [BGV12] and thoroughly studied by Langlois and Stehlé [LS15]. Instead of considering noisy linear equations over the rational integers \mathbb{Z} , the linear equations are now defined algebraically over some ring of polynomials with coefficients in \mathbb{Z} . Intuitively, the improvement in efficiency comes from the fact that every polynomial in this ring defines a structured matrix of multiplication. Thus, we can replace the completely random matrix of LWE by some structured matrix for M-LWE. Additionally, in this case the matrix-vector product can be computed in quasi-linear time. More formally, we now work over some number field K with its associated ring of integers R . For simplicity, we can think of K as the ring of polynomials with coefficients in \mathbb{Q} modulo some irreducible polynomial $f(x)$ of degree n , i.e., $K = \mathbb{Q}[x]/\langle f(x) \rangle$. We call n the degree of K . Thus, elements of K are polynomials of degree less than n with rational coefficients. Its ring of integers can be represented (in most cases) as the polynomials of K whose coefficients are integers, i.e., $R = \mathbb{Z}[x]/\langle f(x) \rangle$. Furthermore, we denote by R_q the quotient ring R/qR , resulting in $R_q = \mathbb{Z}_q[x]/\langle f(x) \rangle$. For positive integers m, d and q , let $\mathbf{A} \in R_q^{m \times d}$ be a matrix sampled uniformly at random over the matrices over R_q . We refer to d as the *rank* of the M-LWE instance. Further, we sample a vector $\mathbf{e} \in R^m$ from some distribution ψ_e over R which has with high probability a small norm (with respect to q) and sample a vector \mathbf{s} from some second distribution ψ_s over R_q^d .¹⁰ Again, we keep both \mathbf{e} and \mathbf{s} secret. An instance of M-LWE is given by $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q) \in R_q^{m \times d} \times R_q^m$. As for the plain LWE problem, it comes in a search variant (requiring to find \mathbf{s}) and a decision variant (asking to distinguish such samples from uniform random ones), depicted in Figure 4.

Why is this variant more efficient than plain LWE? We can interpret the matrix \mathbf{A} over R_q as a matrix over \mathbb{Z} . To do so, we observe that the multiplication of a ring element $a \in R_q$ with another ring element $s \in R_q$ (where both are polynomials with degree less than n) corresponds to the convolution of two polynomials modulo the defining polynomial $f(x)$. The convolution modulo $f(x)$ can be expressed as the matrix-vector product using the matrix of multiplication modulo $f(x)$ of a , denoted by $\text{Rot}(a) \in \mathbb{Z}_q^{n \times n}$. In other words $a \cdot s$ over R_q becomes $\text{Rot}(a) \cdot s$ over \mathbb{Z}_q . Note that the full matrix $\text{Rot}(a)$ is determined by its first column. For simplicity, one can think of the rotation matrix as a nega-cyclic matrix, as illustrated in Figure 4. Hence,

¹⁰For simplicity, we keep the introduction in the *primal* version of M-LWE, where \mathbf{s} is a vector over R . Later, we consider the *dual* version, where the secret \mathbf{s} is a vector over the dual R^\vee . Furthermore, we present the problem in its *discrete* form, where the noise is sampled from a distribution ψ_e over R . Later, we also study the *continuous* version, where ψ_e is a distribution over (an extension of) K .




Figure 4: An instance of the M-LWE problem in its search variant. The number of rows m of \mathbf{A} represents the number of given M-LWE samples and the number of columns d of \mathbf{A} defines the rank of the M-LWE instance. Every entry $a_{jk} \in R_q$ of the matrix \mathbf{A} defines a matrix of multiplication $\text{Rot}(a_{jk}) \in \mathbb{Z}_q^{n \times n}$, where n is the ring degree of R . For simplicity, we illustrate the rotation matrix as a nega-cyclic matrix.

the $m \times d$ matrix $\mathbf{A} = (a_{jk})_{jk} \in R_q^{m \times d}$ defines an $mn \times dn$ matrix $\text{Rot}(\mathbf{A}) = (\text{Rot}(a_{jk}))_{jk}$ over \mathbb{Z}_q . This said, an instance of M-LWE with parameters n and d can be viewed as an instance of LWE of dimension $N = nd$, where the public matrix $\text{Rot}(\mathbf{A})$ is not fully random, but has some structure. This is why we often refer to it as a *structured* variant of LWE.

One interesting special case is M-LWE with rank $d = 1$, which is commonly referred to as the Ring Learning With Errors (R-LWE) problem, originally introduced by Lyubashevsky et al. [LPR10, LPR13]. Recently, Peikert and Pepin [PP19] showed a tight reduction from R-LWE over a number field of degree $n \cdot k$ to M-LWE over a number field of degree n and with rank k . The R-LWE problem is itself closely related to yet another structured variant, the so-called Polynomial Learning With Errors (P-LWE) problem [SSTX09]. For specific choices of number fields K and defining polynomials $f(x)$ both problems are equivalent, as shown by Roçca et al. [RSW18]. The special case of M-LWE, where the ring has degree $n = 1$ and thus $R = \mathbb{Z}$, is exactly the original LWE problem. This is why M-LWE can be seen as a generalization of the unstructured LWE and the structured R-LWE (and hence P-LWE) problems.

In many respects, the more recent problem M-LWE resembles its unstructured and more studied counterpart LWE. For instance, for suitable parameter choices M-LWE also enjoys worst-case to average-case connections from lattice problems such as SVP_γ or GapSVP_γ [LS15]. Whereas the hardness results for LWE start from the lattice problem in the class of general Euclidean lattices, the set has to be restricted to *module lattices* in the case of M-LWE. Such module lattices correspond to modules in the ring R and we refer to the related lattice problems as Mod-SVP_γ , Mod-GapSVP_γ and Mod-SIVP_γ , respectively. Note that we are now considering lattice problems in a specific class of lattices which have an additional algebraic structure. One may rise the question if those lattice problems are still hard when we restrict them to module lattices. In the special case of module lattices of rank 1, i.e., ideal lattices, there are indeed some weaknesses. First, the decision problem Mod-GapSVP_γ becomes easy to solve for ideal lattices, as their minimum can be bounded above *and* below [PR07]. Second, for some classes of number fields the SVP_γ problem is shown to be easier on ideal lattices than on general lattice [CDPR16, CDW17, PHS19, BR20]. However, no algorithm is known so far that solves problems over module lattices of rank strictly larger than 1 faster than on general lattices, lead-

ing to the following conjecture:

🚩 **Conjecture 2:** There is no polynomial-time classical or quantum algorithm that approximates the lattice problems Mod-SVP_γ , Mod-GapSVP_γ or Mod-SIVP_γ to within polynomial factors γ for module lattices of rank at least 2.

Nevertheless, some important properties that have been proven for LWE, have not (yet) been demonstrated for M-LWE. For instance, LWE possesses a classical reduction from worst-case lattice problems as shown by Peikert [Pei09] and later improved by Brakerski et al. [BLP⁺13], which is not known for M-LWE. Furthermore, LWE doesn't become significantly easier, even if the uniform secret distribution ψ_s is replaced by some narrower distribution, yielding for instance binary secrets [GKPV10, BLP⁺13], again, no such result is proven for the module setting. Since its introduction in 2012, M-LWE has enjoyed more and more popularity as it offers a fine-grained trade-off between concrete security and efficiency. Within the NIST standardization process [NIS], several third round candidates rely on the hardness of M-LWE, e.g., the signature scheme Dilithium [DKL⁺18] and the key encapsulation mechanism Kyber [BDK⁺18].

Gaining in efficiency on the positive side comes with a potential decrease in the security level guarantees on the negative side. There are concrete examples of polynomials $f(x)$ for which P-LWE becomes computationally easy: for instance when $f(x)$ 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 P-LWE and R-LWE is given by Peikert [Pei16b]. To the best of our knowledge, it is still not fully understood how to choose a *good* polynomial for instantiating structured variants of LWE. Motivated by the question of how to choose a good polynomial, Roşca et al. [RSSS17] introduce the Middle-Product Learning With Errors (MP-LWE) problem, a variant of LWE, whose hardness does not depend on a *single* polynomial, but on a set of *exponentially* many polynomials. Informally, the middle-product of two polynomials is determined by the middle coefficients of their convolution product, and thus it is independent of some defining polynomial. As for the hardness of MP-LWE, Roşca et al. [RSSS17] establish a reduction from the P-LWE problem parametrized by a polynomial $f(x)$ to the MP-LWE problem defined independently of any such $f(x)$. Thus, as long as the P-LWE problem defined over some f (belonging to a huge family of polynomials) is hard, the MP-LWE problem is also guaranteed to be hard. A more recent result by Peikert and Pepin [PP19] shows a tighter (and more direct) reduction from R-LWE to MP-LWE. As a cryptographic application, Roşca et al. [RSSS17] propose an encryption scheme that is proven secure under the MP-LWE hardness assumption, with keys of size $\tilde{O}(\lambda)$ and running time $\tilde{O}(\lambda)$, where λ is the security parameter.

Contributions

The research that is presented in this manuscript focuses on structured variants of the LWE problem and their use for quantumly secure cryptography. In the following, we outline the shown contributions which can be divided into two different categories. The first part contains contributions to the better understanding of the theoretical foundations of lattice-based cryptography. The second part includes the design of two efficient encryption schemes, whose security proofs are based on structured lattice problems. We illustrate the contributions and contextualize them with respect to existing results in Figure 5.

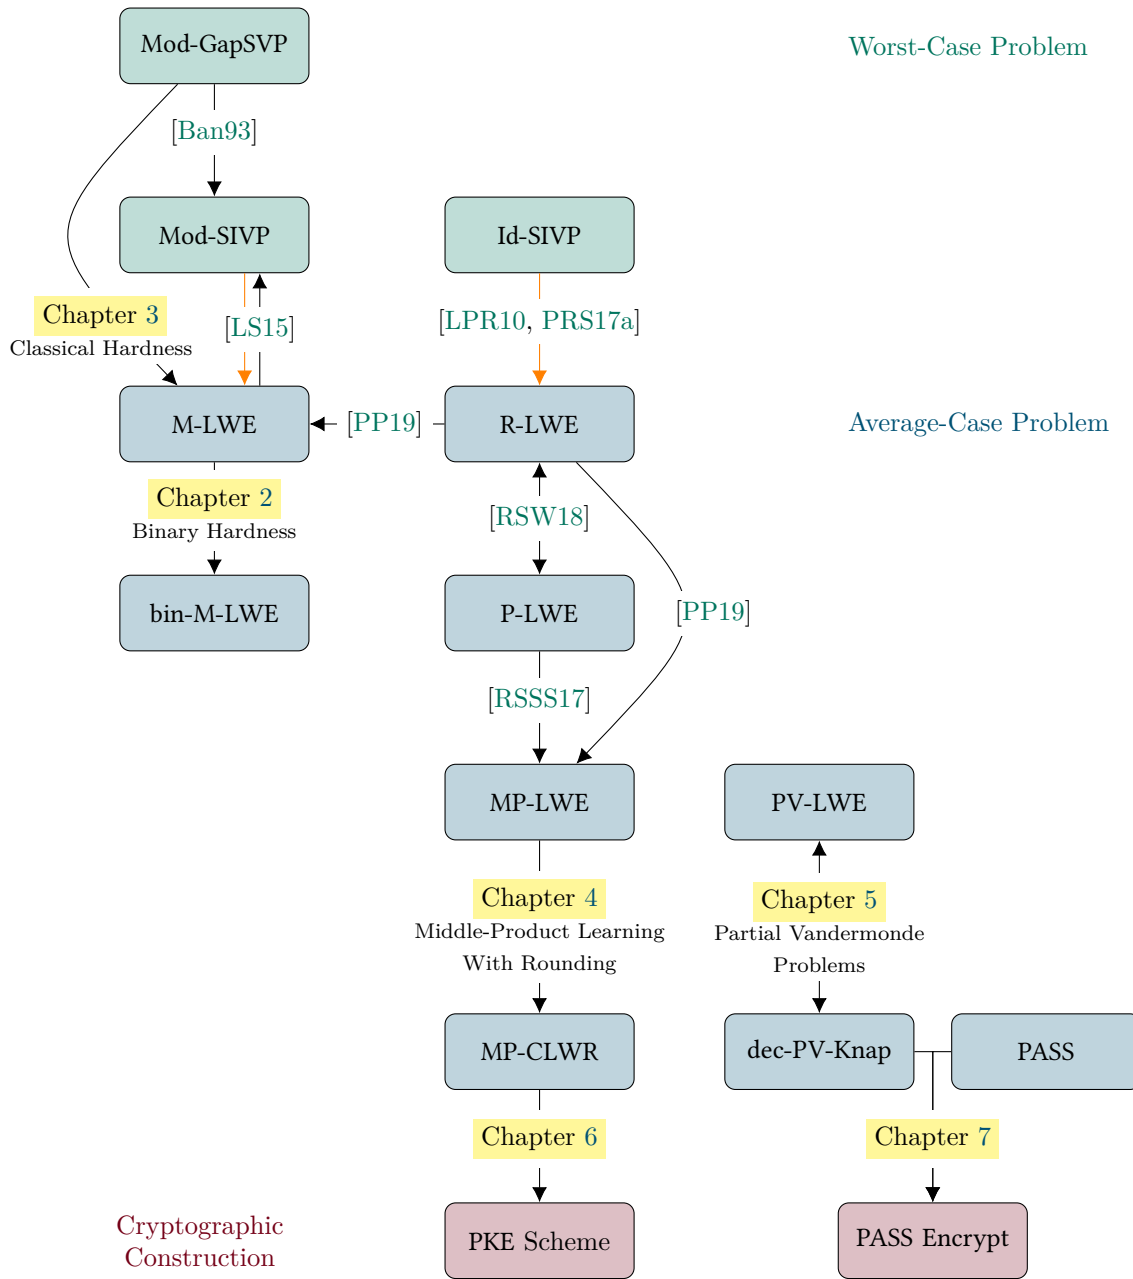


Figure 5: Overview of the contributions presented in this thesis and their connection to existing results. The boxes in green represent worst-case problems over structured lattices, the blue ones are average-case lattice problems and the red boxes design cryptographic constructions. An arrow from one to another box denotes the existence of an efficient reduction from the first to the second problem. In the case of an arrow from a problem to a cryptographic scheme, it says that the security of the scheme is based on the corresponding problem. To ease readability, we omit the associated parameters of the problems, which may be affected by the reduction. An orange arrow means that the reduction is not classical, but quantum.

Theoretical Foundations


As explained before, there are some important characteristics that have been proven for LWE, which have not (yet) been demonstrated for M-LWE. We think that a rigorous study of the M-LWE problem is essential to deepen our confidence in the difficulty assumptions that are made in currently proposed cryptographic schemes for standardization. In the following, we present how we fill two of such knowledge gaps concerning M-LWE.

Binary Hardness

As a first contribution, we prove in Chapter 2 that M-LWE remains difficult to solve, even if the secret distribution ψ_s yielding the vector s over the ring R is changed. Instead of using the uniform distribution which leads to polynomials with large coefficients (i.e., in \mathbb{Z}_q), we can also use a distribution which yields polynomials with binary coefficients (i.e., in $\{0, 1\}$). We denote the latter problem by **bin-M-LWE**. It is particularly interesting as it increases efficiency, enables efficient modulus-rank switching techniques [BLP⁺13, AD17a, WW19] and is needed in Fully Homomorphic Encryption schemes as in [DM15, CGGI16]. For the unstructured case, reductions from LWE to LWE with a binary secret and slightly increased dimension and error bound have been shown [GKPV10, BLP⁺13, Mic18]. However, all existing results work over the integers \mathbb{Z} and not over the ring of integers R of some number field K . Hence, as stated by Micciancio in the conclusion of [Mic18], an important open problem is whether similar results carry over to the structured variants, in particular to M-LWE. In this thesis we make progress towards solving this problem by proving the hardness of **bin-M-LWE**, if the module rank d is (super-)logarithmic in the degree n of the underlying number field K . More precisely, we present two different proofs to obtain this result, which slightly differ in their requirements, achieved parameters and concrete variants of the problem.

Moreover, we show that both reductions can be generalized to a larger secret distribution ψ_s , where the secret vector s over the ring R has coefficients in $\{0, \dots, \eta-1\}$ for some positive integer η . Thus, **bin-M-LWE** corresponds to the special case of $\eta = 2$. Using secrets with larger coefficients increases the resulting error linearly in η , but weakens the rank condition by a logarithmic factor.

The results of this chapter imply the hardness of M-LWE with a small (with respect to its coefficients) secret and a moderate rank (e.g., $\Omega(\log_2 n)$), which is guaranteed even in the presence of *arbitrarily* many samples. For a flexible choice of parameters, current NIST candidates as **Dilithium** [DKL⁺18] and **Kyber** [BDK⁺18] consider M-LWE variants with a small secret and also a small rank, while restricting the number of samples to be small (e.g., linear in n) to rule out the BKW type of attacks [KF15].


 **Reference:** The content of this chapter, besides the generalization to larger secret distributions, is based on two joint works with Corentin Jeudy, Adeline Roux-Langlois and Weiqiang Wen. The first is published in the proceedings of the conference Asiacrypt 2020 [BJRW20] and the second is published in the proceedings of the conference CT-RSA 2021 [BJRW21].

Classical Hardness

As a second contribution concerning the theoretical hardness of M-LWE, we strengthen its connection to worst-case problems over structured Euclidean lattices. More precisely, we prove in Chapter 3 a classical reduction from the worst-case problem **Mod-GapSVP** over module lattices to M-LWE. As a result, a classical and efficient algorithm that solves M-LWE could be used to

solve the decision version of the Shortest Vector Problem over module lattices, which is supposed to be intractable. Prior to this, only a quantum reduction from the related problem Mod-SIVP over module lattices was known [LS15].¹¹ Such a reduction requires quantum computers, which are extremely powerful, expensive to construct and not available on a large-scale yet. In contrast, the reduction we present in this thesis only needs standard computers. For the unstructured case, a classical hardness proof from GapSVP to LWE is given by Brakerski et al. [BLP⁺13].

At a high level, we follow the same structure as in [BLP⁺13] by rigorously adapting all needed results to the module setting. In more details, we need three ingredients: First, we lack a classical reduction from Mod-GapSVP to decision M-LWE with an exponential-sized modulus, which we obtain by adapting the LWE counterpart by Peikert [Pei09] that we combine with a more recent result by Peikert et al. [PRS17a]. As a second component, we need the hardness of M-LWE using a binary secret, which we have proven in Chapter 2. Finally, a modulus reduction technique is required, where we can use the modulus-rank switching technique shown by Albrecht and Deo [AD17a]. By carefully putting together all three ingredients we prove a classical reduction from Mod-GapSVP with module rank at least 2 to M-LWE for any polynomial-sized modulus p and module rank d at least $3n + \omega(\log_2 n)$, where n is the degree of the underlying number field K .

 **Reference:** This chapter can be seen as a continuation of Chapter 2 and is therefore based on the same joint work with Corentin Jeudy, Adeline Roux-Langlois and Weiqiang Wen [BJRW20]. By using the improved results for M-LWE with a binary secret from [BJRW21], we obtain a simpler and tighter proof compared to the original version in the proceedings of the conference Asiacrypt 2020.

Middle-Product Learning With Rounding


After having studied variants of the M-LWE problem which serve as the underlying hardness assumptions of current practical encryption schemes and digital signatures, we turn in Chapter 4 our attention to a completely new structured version of LWE. The main motivation of this chapter is to combine two flavors of LWE that we have presented before, the middle-product variant MP-LWE and the rounding variant LWR, in order to define a new problem that benefits from both of their advantages. To this end, we introduce a new hardness assumption which we refer to as the Middle-Product Computational Learning With Rounding (MP-CLWR) problem. On the one hand, MP-CLWR uses rounding in a similar way to LWR and hence avoids the Gaussian error sampling. On the other hand, the hardness of MP-CLWR does not depend on a specific defining polynomial $f(x)$, but is guaranteed by the hardness of P-LWE for an exponentially large set of defining polynomials. Thus, the MP-CLWR assumption enjoys the desired properties from both, the security advantage of MP-LWE and the simplicity advantage of LWR.

The reason why we introduce the problem in its *computational* form, instead of the more standard *search* or *decision* variant, is twofold: On the one hand, it is as for today unclear how to reduce the hardness of decision MP-LWR from worst-case lattice problems, while maintaining the coefficient-wise rounding and allowing for a polynomially large modulus. On the other hand, it is unclear how to construct encryption schemes directly on the search variant. The computational notion solves this issue, as we can derive an efficient reduction still allowing for a polynomially large modulus and at the same time we can use it to build an encryption scheme (which we introduce separately in Chapter 6).

Regarding the difficulty of this new problem, we prove that MP-CLWR is at least as hard

¹¹As illustrated in Figure 5, Mod-SIVP is at least as hard as Mod-GapSVP using Banaszczyk’s transference theorem [Ban93].

as MP-LWE, which is itself based on the hardness of P-LWE for some defining polynomial $f(x)$ which belongs to an exponentially large set.

 **Reference:** The results of this chapter are based on a joint work with Shi Bai, Dipayan Das, Adeline Roux-Langlois, Weiqiang Wen and Zhenfei Zhang which is published in the proceedings of the conference Asiacrypt 2019 [BBD⁺19].

Partial Vandermonde Problems


In the chapters before, we study flavors of LWE whose hardness is implied by the presumed intractability of worst-case problems over structured Euclidean lattices. We can see both, the module variants of LWE that we investigate in Chapter 2 and Chapter 3, and the middle-product version of LWR from Chapter 4, as descendants of Ajtai’s [Ajt96] and Regev’s [Reg05] works. As mentioned before, at the end of the 1990s a second line of work was initiated by Hoffstein et al. [HPS98], better known as the NTRU encryption scheme. NTRU-derived schemes are usually characterized by their efficiency and simplicity, but lack a connection to worst-case lattice problems.¹² Note that among the three lattice-based finalists of NIST’s standardization process [NIS], there is also a variant of the original NTRU scheme [CDH⁺].

In Chapter 5, we study another family of average-case problems on structured lattices which is related to the discrete Vandermonde transform and that, like the NTRU problem, allows for simple and efficient protocols at the expense of having no known worst-case to average-case connection. In more details, we consider a cyclotomic field K of degree n with associated ring of integers R . A cyclotomic field is a number field whose defining polynomial $f(x)$ is a cyclotomic polynomial (having primitive roots of unity as its roots). Further, we consider a prime q such that $f(x)$ completely splits over \mathbb{Z}_q , thus, we require $f(x)$ to possess exactly n distinct roots in the field \mathbb{Z}_q . The discrete Vandermonde transform $\bar{V} \cdot a$ of a polynomial $a \in R$ is the evaluation of this polynomial at the n roots of $f(x)$ over \mathbb{Z}_q . The conjecture, building the basis of this chapter, is that there is no polynomial-time classical or quantum algorithm that recovers an element of R with small coefficients by having access only to a partial list of its discrete Vandermonde transform. Of course, it is crucial to define how *small* the coefficients of the polynomial have to be and how *large* the partial list is that we provide. In analogy to the standard Knapsack problem, we call this the Partial Vandermonde Knapsack problem (PV-Knap). A related problem is Partial Vandermonde SIS (PV-SIS), where, given a subset of the roots of the defining polynomial $f(x)$, one asks to find a ring element of small norm whose evaluation at those roots is zero in \mathbb{Z}_q . Again, PV-SIS can be seen as specific variant of the standard Short Integer Solution problem, and hence it can be formulated as a Shortest Vector Problem over a structured lattice.

In 2014, Hoffstein et al. [HPS⁺14] proposed a digital signature scheme called PASS Sign, whose provable security is based on the difficulty of solving PV-Knap and PV-SIS [LZA18]. Shortly afterwards, Hoffstein and Silverman [HS15] introduce PASS Encrypt, a Public Key Encryption (PKE) scheme whose computational building blocks are closely related to the ones of PASS Sign. Their encryption scheme is very efficient and fulfills additive and multiplicative homomorphic properties, which make it a natural starting point for the design of efficient cryptographic primitives. For example, such properties are recently exploited in the context of PASS Sign to construct compact aggregate signature schemes [DHSS20], and it is plausible that combining PASS Encrypt with PASS Sign may form the basis for various compact and efficient privacy-preserving primitives such as group signatures. Unfortunately, a main drawback of PASS Encrypt to date is that

¹²There are variants of NTRU with security proofs from worst-case problems over structured Euclidean lattices [SS11, SS13, YXW17], but they are rarely used in practice for efficiency reasons.

no proof of security was given in [HS15] with respect to the hardness of explicit computational problems. In this chapter, we make progress towards understanding the hardness assumptions needed to prove the security of PASS Encrypt. First, we emphasize its connection to average-case ideal lattices, even though we can't show a worst-case to average-case reduction as we have for structured variants of LWE. Second, we enlarge the landscape of problems that use the partial Vandermonde transform by defining a new variant of LWE, called Partial Vandermonde Learning With Errors (PV-LWE). Later, we show the equivalence of PV-Knap and PV-LWE by exploiting the same duality connection as we have for standard Knapsack problems and LWE. As our main motivation is to provide a security proof for PASS Encrypt, we define a variant of PV-Knap, that we call the PASS problem. This problem serves (together with the decision version of PV-Knap) as the underlying hardness assumption for (a slightly modified version of) PASS Encrypt. We present the scheme together with the security proof later in Chapter 7. We conclude the chapter by showing that for a special choice of the partial Vandermonde transform, one can accelerate computations when working with PV-Knap and PV-LWE.


 **Reference:** The contributions presented in this chapter are based on a joint work with Amin Sakzad and Ron Steinfeld, which was initiated during a research stay at the Monash University in Melbourne from October to December 2019. The results haven't been published before and are publicly presented for the first time in this thesis.

Cryptographic Constructions

In the first part of the manuscript, we introduce two new hardness assumptions related to problems over structured Euclidean lattices. In order to provide insights in how to use them for cryptography, we construct two Public Key Encryption (PKE) schemes in the second part of this work. Both are proven to be correct and secure assuming explicitly stated computational assumptions. To this end, we use for both schemes the standard notion of IND-CPA security. Informally speaking, this security notion captures that no efficient adversary can distinguish between the ciphertext of two messages, where we allow the adversary to choose the messages by themselves.


Encryption Based on Middle-Product LWR

To show the cryptographic usefulness of our newly defined hardness assumption MP-CLWR from Chapter 4, we present in Chapter 6 a PKE scheme whose hardness is implied by MP-CLWR. Its design is simultaneously inspired by two existing PKE schemes: On the one hand, the PKE scheme from Roşca et al. [RSSS17] whose security is based on MP-LWE, and on the other hand, the PKE scheme from Chen et al. [CZZ18] whose security is based on the hardness of the computational variant of LWR over rings. As for [CZZ18], we make use of a so-called reconciliation mechanism, which guarantees the correctness of our scheme. We prove our scheme's correctness and security based on the hardness of MP-CLWR in the idealized Random Oracle Model. As typical for LWR-based schemes, during public key generation we only need to round the middle-product of two polynomials instead of sampling Gaussian error, making it easier to implement than the scheme in [RSSS17]. Simultaneously, while guaranteeing an equal level of security, we obtain the same asymptotic key and ciphertext sizes as the PKE scheme of [RSSS17]. Furthermore, we save in bandwidth, as the second part of the public key is modulo p and not modulo q , where $p \ll q$ by some order of magnitude. Finally, we analyze the concrete security of our scheme by looking at the currently best known attacks against it.

 **Reference:** This chapter can be seen as the continuation of Chapter 4 and is therefore based on the same joint work with Shi Bai, Dipayan Das, Adeline Roux-Langlois, Weiqiang Wen and Zhenfei Zhang [BBD⁺19].

PASS Encrypt

As our final contribution, we present in Chapter 7 a modified version of PASS Encrypt together with a security proof based on the decision PV-Knap problem and a leaky variant of it, that we call the PASS problem, both studied in Chapter 5. The latter problem captures the fact that a ciphertext of PASS Encrypt consists of several partial Vandermonde transforms of *related* elements. Hence, a successful attacker against PV-Knap can be used to win the IND-CPA security game of PASS Encrypt, but a successful attacker against the IND-CPA security may not be powerful enough to solve PV-Knap. This issue was not addressed before in the original version of PASS Encrypt by Hoffstein and Silverman [HS15]. Furthermore, in the original version, the scheme is deterministic and thus cannot be IND-CPA secure. To solve this issue, we make the scheme probabilistic by adding randomness to the message. We then give a proof of correctness for well-chosen parameters and a proof of security, assuming the hardness of decision PV-Knap and PASS. We also provide a refined analysis of the scheme's security, where we show a novel attack that we call Plaintext Recovering Using Hints attack, which takes the structure of PASS Encrypt into account. To this end, we make use of a recent work by Dachman-Soled et al. [DDGR20] to analyze instances of LWE, where additional hints on the secret and/or error are given. Our complexity estimates for this attack show that it does not reduce the attack complexity below that of previously known lattice attacks on PASS, which increases our confidence in the claimed security of PASS Encrypt against best known lattice attacks. We conclude the chapter by providing concrete sample parameters and compare our scheme with two other efficient schemes whose security proofs are based on structured lattice problems.

 **Reference:** This chapter can be seen as the continuation of Chapter 5 and is therefore based on the same joint work with Amin Sakzad and Ron Steinfeld, which was initiated during a research stay at the Monash University in Melbourne from October to December 2019.

Chapter 1

Preliminaries

In this chapter, we recall the mathematical and cryptographic notions needed within this thesis. This encompasses the presentation of the necessary objects of algebraic number theory, lattice theory and probability theory. We further introduce the Learning With Errors problem in its various flavors and conclude with the cryptographic definitions used within this work.

Contents

1.1 Algebraic Number Theory	2
1.1.1 Space H	3
1.1.2 Canonical Embedding, Trace and Norm	3
1.1.3 Coefficient Embedding and Rotation Matrix	4
1.1.4 Discrete Vandermonde Matrix	5
1.1.5 Modules, Ideals and Units	6
1.2 Lattice Theory	7
1.2.1 Euclidean Lattices	7
1.2.2 Structured Lattices	8
1.2.3 Computational Problems	9
1.3 Probability Theory	10
1.3.1 Gaussian Measures	10
1.3.2 Gaussians over Number Fields	13
1.3.3 Statistical Distance	16
1.3.4 Rényi Divergence	17
1.3.5 Leftover Hash Lemma	18
1.4 Learning With Errors	19
1.4.1 Learning With Errors (LWE)	19
1.4.2 Polynomial Learning With Errors (P-LWE)	20
1.4.3 Module Learning With Errors (M-LWE)	21
1.4.4 Middle-Product Learning With Errors (MP-LWE)	23
1.4.5 Learning With Rounding (LWR)	25
1.4.6 Practical Hardness of Learning With Errors	27
1.5 Cryptographic Notions	27
1.5.1 Public Key Encryption	27
1.5.2 Random Oracle Model	28

Within this thesis, we use the standard Landau notations (e.g., $\omega(\cdot)$, $\Omega(\cdot)$, $O(\cdot)$ and $\Theta(\cdot)$) and we call a function ε negligible in n if $\varepsilon(n) = n^{-\omega(1)}$, i.e., it decreases faster towards 0 than the inverse of any polynomial function. Sometimes, we simply write $\text{negl}(n)$ to denote a negligible function in n . The abbreviation PPT stands for probabilistic polynomial-time.

Throughout the manuscript, \mathbb{N} denotes the set of natural numbers without 0. For any $q \in \mathbb{N}$, we denote by \mathbb{Z}_q the integers modulo q . For simplicity, we denote by $[n]$ the set $\{1, \dots, n\}$ for any $n \in \mathbb{N}$. We denote the Kronecker symbol by $\delta_{j,k}$ which equals 1 if $j = k$ and 0 otherwise. Column vectors are written in bold lowercase letters \mathbf{a} and matrices in bold uppercase letters \mathbf{A} . The transpose and Hermitian operators over matrices are denoted by \mathbf{A}^T and \mathbf{A}^\dagger . The concatenation of two matrices \mathbf{A} and \mathbf{B} is denoted by $[\mathbf{A}|\mathbf{B}]$ and their Kronecker product by $\mathbf{A} \otimes \mathbf{B}$. The canonical basis of \mathbb{C}^n is given by $\{\mathbf{e}_j\}_{j \in [n]}$, where \mathbf{e}_j has 1 as its j -th coefficient and else only 0's. For a vector $\mathbf{a} \in \mathbb{C}^n$, we define the matrix $\text{diag}(\mathbf{a})$ to be the diagonal matrix whose diagonal entries are the entries of \mathbf{a} . For any vector $\mathbf{a} = (a_1, \dots, a_n)^T$, we denote by $\text{rev}(\mathbf{a}) = (a_n, \dots, a_1)^T$ the vector in reverse order. The identity matrix of size $n \times n$ is denoted by \mathbf{I}_n . For any $\mathbf{a} \in \mathbb{C}^n$, we define the Euclidean norm as $\|\mathbf{a}\|_2 = \sqrt{\sum_{j \in [n]} |a_j|^2}$ and the infinity norm as $\|\mathbf{a}\|_\infty = \max_{j \in [n]} |a_j|$. We also define the spectral norm of any matrix $\mathbf{A} = (a_{jk})_{j \in [n], k \in [m]} \in \mathbb{C}^{n \times m}$ as $\|\mathbf{A}\|_2 = \max_{\mathbf{x} \in \mathbb{C}^m \setminus \{0\}} \|\mathbf{A}\mathbf{x}\|_2 / \|\mathbf{x}\|_2$, and the max norm as $\|\mathbf{A}\|_{\max} = \max_{j \in [n], k \in [m]} |a_{j,k}|$. The Gram-Schmidt orthogonalization of a matrix \mathbf{A} from left to right is denoted by $\text{GS}(\mathbf{A})$.

For a complex number z , we denote by $\Re(z)$ its real component and \bar{z} its complex conjugate. The singular values of a matrix \mathbf{A} are the square roots of the non-negative eigenvalues of $\mathbf{A}^\dagger \mathbf{A}$ and we denote its largest singular value by $\mathfrak{s}_1(\mathbf{A})$. In particular, the singular values of a diagonal matrix are given by the absolute values of its diagonal entries. It can be shown that the spectral norm $\|\mathbf{A}\|_2$ of a matrix corresponds to its largest singular value.

1.1 Algebraic Number Theory

In this section, we give a brief review of the number-theoretic notions that are used in the context of lattice-based cryptography. For further details, we refer to the works of Lyubashevsky et al. [LPR10, LPR13] and the short note of Conrad [Con].

An algebraic number $\zeta \in \mathbb{C}$ is a complex root of a polynomial over \mathbb{Q} . Its minimal polynomial $f(x) \in \mathbb{Q}[x]$ is the unique irreducible monic polynomial of minimal degree such that ζ is one of its roots. Sometimes, we also call it the defining polynomial. An algebraic number is called an algebraic integer if its minimal polynomial has only integer coefficients, i.e., $f(x) \in \mathbb{Z}[x]$. A number field $K = \mathbb{Q}(\zeta)$ of degree n is a finite extension of the rational number field \mathbb{Q} obtained by adjoining an algebraic number ζ . The degree n of K is given by the degree of the minimal polynomial of ζ . We define the tensor field $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$ which can be seen as the finite field extension of the reals by adjoining ζ . The set of all algebraic integers of K defines a ring, called the ring of integers, which we denote by R . It is always true that $\mathbb{Z}[\zeta] \subseteq R$, where this inclusion can be strict.

Some results of this thesis are restricted to the class of number fields, where the equality $R = \mathbb{Z}[\zeta]$ holds. This is the case for some quadratic extensions (i.e., when $\zeta = \sqrt{d}$ with d square-free and $d \not\equiv 1 \pmod{4}$), cyclotomic fields (i.e., when ζ is a primitive root of unity) and number fields with a defining polynomial $f(x)$ of square-free discriminant Δ_f . The class of cyclotomic number fields plays an important role in this manuscript. For $\nu \in \mathbb{N}$ let $\zeta \in \mathbb{C}$ be a ν -th root of unity, i.e., $\zeta^\nu = 1$. Then, the ν -th cyclotomic number field $K = \mathbb{Q}(\zeta)$ has degree $n = \varphi(\nu)$, where φ is Euler's totient function. In this case, the minimal polynomial is $f(x) = \prod_{j \in [n]} (x - \alpha_j)$, where the α_j are the distinct primitive ν -th roots of unity.

Example 1.1 (Power-of-2 Cyclotomic Fields)

One particularly interesting subclass of cyclotomic fields, from a theoretical and practical point of view, is given by cyclotomic fields where ν is a power of 2, i.e., $\nu = 2^{\ell+1}$ for some $\ell \in \mathbb{N}$. In this case, it yields $n = \varphi(\nu) = 2^\ell$ and $f(x) = x^n + 1$. In the following, we take this as our running example to illustrate properties and constructions defined in this section. Furthermore, several results of this thesis are restricted to this special class of number fields, as for instance Section 3.4.2, Section 5.3 and Section 5.4.

1.1.1 Space H

Let $K = \mathbb{Q}(\zeta)$ be a number field of degree n . By t_1 we denote the number of real roots of the minimal polynomial of ζ , and by t_2 the number of pairs of complex conjugate roots, such that $n = t_1 + 2t_2$. We introduce the space $H \subseteq \mathbb{R}^{t_1} \times \mathbb{C}^{2t_2}$ as

$$H = \left\{ \mathbf{x} = (x_1, \dots, x_n)^T \in \mathbb{R}^{t_1} \times \mathbb{C}^{2t_2} : x_{t_1+t_2+j} = \overline{x_{t_1+j}}, \forall j \in [t_2] \right\}.$$

For $j \in [t_1]$, we set $\mathbf{h}_j = \mathbf{e}_j$, and for $j \in \{t_1 + 1, \dots, t_1 + t_2\}$, we set $\mathbf{h}_j = \frac{1}{\sqrt{2}}(\mathbf{e}_j + \mathbf{e}_{j+t_2})$ and $\mathbf{h}_{j+t_2} = \frac{i}{\sqrt{2}}(\mathbf{e}_j - \mathbf{e}_{j+t_2})$, where i denotes the complex imaginary number $i = \sqrt{-1}$. The set $\{\mathbf{h}_j\}_{j \in [n]}$ forms an orthonormal basis of H as a real vector space, showing that H is isomorphic to \mathbb{R}^n . The change of basis is given by the unitary matrix

$$\mathbf{H} = \begin{bmatrix} \mathbf{I}_{t_1} & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}}\mathbf{I}_{t_2} & \frac{i}{\sqrt{2}}\mathbf{I}_{t_2} \\ 0 & \frac{1}{\sqrt{2}}\mathbf{I}_{t_2} & \frac{-i}{\sqrt{2}}\mathbf{I}_{t_2} \end{bmatrix},$$

where \mathbf{I}_n denotes the $n \times n$ identity matrix.

1.1.2 Canonical Embedding, Trace and Norm

Any number field $K = \mathbb{Q}(\zeta)$ of degree n has exactly n field embeddings (i.e., injective field homomorphisms) $\sigma_j: K \rightarrow \mathbb{C}$ fixing each element of \mathbb{Q} , where $j \in [n]$. Let $\sigma_1, \dots, \sigma_{t_1}$ be the real embeddings (i.e., the embeddings whose image lies in \mathbb{R}) and $\sigma_{t_1+1}, \dots, \sigma_{t_1+2t_2}$ the complex embeddings (i.e., the embeddings whose image lies in \mathbb{C}). The complex ones come in conjugate pairs, thus $\sigma_j = \overline{\sigma_{j+t_2}}$ for $j \in \{t_1 + 1, \dots, t_1 + t_2\}$. In the case of cyclotomic fields, all n embeddings are complex, thus $t_1 = 0$ and $t_2 = \frac{n}{2}$.

The *canonical embedding* σ is the map $\sigma: K \rightarrow H$, defined by $\sigma(\mathbf{x}) = (\sigma_j(x))_{j \in [n]}$. It describes a field homomorphism, where multiplication and addition in H are component-wise. Hence, $\sigma(a + b) = \sigma(a) + \sigma(b)$ and $\sigma(a \cdot b) = \sigma(a) \circ \sigma(b)$ for $a, b \in K$, where \circ denotes the component-wise product. By abuse of notation, for $d \in \mathbb{N}$ and a vector $\mathbf{x} \in K^d$, we write $\sigma(\mathbf{x})$ to denote the vector $(\sigma(x_k))_{k \in [d]} \in H^d \subseteq \mathbb{C}^{nd}$. We can represent $\sigma(\mathbf{x}) \in H$ via the real vector $\sigma_H(\mathbf{x}) \in \mathbb{R}^n$ through the change of basis described above, i.e., $\sigma_H(\mathbf{x}) = \mathbf{H}^\dagger \cdot \sigma(\mathbf{x})$. Note that, as opposed to σ , multiplication is not component-wise for σ_H . More concretely, in the basis $\{\mathbf{e}_j\}_{j \in [n]}$, multiplication by $x \in K$ can be described as the left multiplication by the diagonal matrix $\mathbf{X} = \text{diag}(\sigma(x))$. Hence, changing to the basis $\{\mathbf{h}_j\}_{j \in [n]}$ leads to the corresponding matrix of multiplication $\mathbf{X}_H = \mathbf{H}^\dagger \cdot \mathbf{X} \cdot \mathbf{H}$, having the same singular values as \mathbf{X} , given by $|\sigma_j(x)|$ for $j \in [n]$.

We use the canonical embedding σ to define two different norms on vectors over K . The first can be seen as the standard infinity norm, after having applied the canonical embedding. The

second is a mixture between the Euclidean norm (with respect to the different coefficients of the vector over K) and the infinity norm (with respect to the canonical embedding). For an element $\mathbf{x} \in K^d$, we define $\|\mathbf{x}\|_\infty = \max_{j \in [n], k \in [d]} |\sigma_j(x_k)|$, and $\|\mathbf{x}\|_{2,\infty} = \max_{j \in [n]} \sqrt{\sum_{k \in [d]} |\sigma_j(x_k)|^2}$.

The trace $\text{Tr}: K \rightarrow \mathbb{Q}$ and the norm $N: K \rightarrow \mathbb{Q}$ are defined as the sum and product, respectively, of the embeddings, i.e., for any $x \in K$, we have $\text{Tr}(x) = \sum_{j \in [n]} \sigma_j(x)$. and $N(x) = \prod_{j \in [n]} \sigma_j(x)$.

1.1.3 Coefficient Embedding and Rotation Matrix

Every number field $K = \mathbb{Q}(\zeta)$ of degree n defines an n -dimensional vector space over \mathbb{Q} with basis $\{1, \zeta, \dots, \zeta^{n-1}\}$. Thus, every element $x \in K$ can be written as $x = \sum_{j=0}^{n-1} x_j \zeta^j$, where $x_j \in \mathbb{Q}$. The *coefficient embedding* $\tau: K \rightarrow \mathbb{Q}^n$ is the isomorphism that sends every element $x \in K$ to its coefficient vector $\tau(x) = (x_0, \dots, x_{n-1})^T$. We also extend the coefficient embedding to $K_{\mathbb{R}}$, which yields an isomorphism between $K_{\mathbb{R}}$ and \mathbb{R}^n . Multiplication by x in the coefficient embedding can be represented by a matrix multiplication, where we denote the corresponding matrix by $\text{Rot}(x) \in \mathbb{R}^{n \times n}$. More concretely, for any $x, y, z \in K$ with $x = yz$, it yields $\tau(x) = \text{Rot}(y) \cdot \tau(z)$. Note that the matrix $\text{Rot}(x)$ is invertible in K for every $x \neq 0$ and that its concrete form depends on the field K .

As both embeddings play an important role in this thesis, we recall how to go from one to the other. For any $x \in K$, the embeddings $\sigma(x)$ and $\tau(x)$ are linked through the Vandermonde matrix \mathbf{V} of the roots of the defining polynomial $f(x)$. For $j \in [n]$, we let $\alpha_j = \sigma_j(\zeta)$ be the j -th root of $f(x)$. Then, $\sigma(x) = \mathbf{V} \cdot \tau(x)$, where

$$\mathbf{V} = \begin{bmatrix} 1 & \alpha_1 & \cdots & \alpha_1^{n-1} \\ 1 & \alpha_2 & \cdots & \alpha_2^{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha_n & \cdots & \alpha_n^{n-1} \end{bmatrix} \in \mathbb{C}^{n \times n}.$$

Example 1.2 (Power-of-2 Cyclotomic Fields)

Looking at the example where $K = \mathbb{Q}(\zeta)$ is the ν -th cyclotomic field with ν a power of 2 and thus $K \cong \mathbb{Q}[x]/\langle x^n + 1 \rangle$ with $n = \nu/2$, the matrix of multiplication is nega-circulant. More precisely,

$$\text{Rot}(x) = \begin{bmatrix} x_0 & -x_{n-1} & \cdots & -x_1 \\ x_1 & x_0 & \ddots & \vdots \\ \vdots & \vdots & \ddots & -x_{n-1} \\ x_{n-1} & x_{n-2} & \cdots & x_0 \end{bmatrix} \in \mathbb{R}^{n \times n}.$$

Using the above, the product of two elements $a, b \in K$ with respect to the coefficient embedding τ fulfills $\|\tau(a \cdot b)\|_\infty = \|\text{Rot}(a) \cdot \tau(b)\|_\infty \leq \|\text{Rot}(a)\|_1 \cdot \|\tau(b)\|_\infty$. Furthermore, the Vandermonde matrix fulfills $\mathbf{V}^\dagger \mathbf{V} = n \cdot \mathbf{I}_n$, thus the singular values of \mathbf{V} all equal \sqrt{n} and hence multiplying by \mathbf{V} is an isometry with scaling factor \sqrt{n} .

⚠ This is not true for more general power-of-prime cyclotomic fields, where the singular values of \mathbf{V} are well known, but are in general not all the same (see [YXW17, Lem. 14]) and thus multiplying by \mathbf{V} is not a scaled isometry.

This transformation is not necessarily structure preserving, e.g., a vector of small norm in the coefficient embedding doesn't need to be of small norm in the canonical embedding as well.

This is captured by the inequalities

$$\|\mathbf{V}^{-1}\|_2^{-1} \|\tau(x)\|_2 \leq \|\sigma(x)\|_2 \leq \|\mathbf{V}\|_2 \|\tau(x)\|_2. \quad (1.1)$$

Hence, the spectral norm of \mathbf{V} and its inverse \mathbf{V}^{-1} help approximating the distortion between both embeddings. Roşca et al. [RSW18] give additional insight on this distortion for specific number fields. Later in Chapter 2, we need for some positive integer η a bound on the parameter $B = \max_{x \in R_\eta} \|\sigma(x)\|_\infty$, where R_η denotes the set of ring elements with a coefficient vector in $\{0, \dots, \eta-1\}^n$. This parameter is inherent to the ring and intervenes in the proof of Lemma 2.2, Lemma 2.3 and Lemma 2.5. Here, we provide an upper-bound on B with respect to the maximum norm of \mathbf{V} , that is further simplified for cyclotomic number fields.

Lemma 1.1 (Upper Bound on B)

For $n \in \mathbb{N}$, let K be a number field of degree n with R its ring of integers. For $\eta \in \mathbb{N}$, we define the set $R_\eta = \{x \in R: \tau(x) \in \{0, \dots, \eta-1\}^n\}$ and let \mathbf{V} be the associated Vandermonde matrix. Then, $B := \max_{x \in R_\eta} \|\sigma(x)\|_\infty \leq n(\eta-1) \|\mathbf{V}\|_{\max}$. In particular, for cyclotomic fields, it yields $B \leq n(\eta-1)$.

Proof: We can express $x \in R_\eta$ as $x = \sum_{j=0}^{n-1} b_j \zeta^j$, with $b_j \in \{0, \dots, \eta-1\}$ for all $0 \leq j \leq n-1$. Then, for $k \in [n]$, we obtain

$$|\sigma_k(x)| \leq \sum_{j=0}^{n-1} b_j |\sigma_k(\zeta)|^j = \sum_{j=0}^{n-1} b_j |\alpha_k|^j \leq \|\mathbf{V}\|_{\max} \sum_{j=0}^{n-1} b_j \leq n(\eta-1) \|\mathbf{V}\|_{\max}.$$

Taking the maximum over all $k \in [n]$ yields $B \leq n(\eta-1) \|\mathbf{V}\|_{\max}$. In the case of cyclotomic fields, the α_k are roots of unity and therefore, all the entries of \mathbf{V} have magnitude 1. Hence $\|\mathbf{V}\|_{\max} = 1$ which yields $B \leq n(\eta-1)$ in this case. ■

1.1.4 Discrete Vandermonde Matrix

In the case of cyclotomic number fields, we can easily control the splitting behavior of the minimal polynomial over the finite field \mathbb{Z}_q , for q prime. More precisely, let $K = \mathbb{Q}[x]/\langle f(x) \rangle$ be the ν -th cyclotomic field with minimal polynomial $f(x)$ of degree $\varphi(\nu) = n$ and let q be prime. If $q \equiv 1 \pmod{\nu}$, then $f(x)$ completely splits in $\mathbb{Z}_q[x]$, i.e., $f(x) = \prod_{j \in [n]} (x - \omega_j) \pmod{q}$, where every ω_j is a distinct primitive ν -th root of unity in \mathbb{Z}_q . Simultaneously, the ideal generated by q over the ring of integers $R = \mathbb{Z}[x]/\langle f(x) \rangle$ has exactly n prime ideal factors defined by the ω_j , i.e., $\langle q \rangle = \prod_{j \in [n]} \langle q, x - \omega_j \rangle$. In this case, the ν -th cyclotomic number field $K = \mathbb{Q}(\zeta)$ with ζ a complex primitive ν -th root of unity, possesses exactly n field homomorphisms $\bar{\sigma}_j: K \rightarrow \mathbb{Z}_q$ for $j \in [n]$, that map ζ to each of the distinct primitive roots ω_j over \mathbb{Z}_q . The *discrete canonical embedding*¹ $\bar{\sigma}$ is the field homomorphism from K to \mathbb{Z}_q^n , defined as $\bar{\sigma}(a) = (\bar{\sigma}_j(a))_{j \in [n]}$. As for the canonical embedding σ over the complex numbers, as introduced in Section 1.1.2, the addition and multiplication are component-wise with respect to $\bar{\sigma}$. Hence, $\bar{\sigma}(a+b) = \bar{\sigma}(a) + \bar{\sigma}(b)$ and $\bar{\sigma}(a \cdot b) = \bar{\sigma}(a) \circ \bar{\sigma}(b)$ for $a, b \in K$. Furthermore, we can again link the coefficient embedding τ and the discrete

¹As the map is not injective, it is not an embedding in its strict sense. We think, however, that its close resemblance to the canonical embedding justifies the word embedding here. In other works, this is also called the Number Theoretic Transform (NTT).

canonical embedding $\bar{\sigma}$ via the Vandermonde matrix of the roots of $f(x)$ in \mathbb{Z}_q . To distinguish this Vandermonde matrix from the one over \mathbb{C} that we defined above, we denote the discrete one over \mathbb{Z}_q by $\bar{\mathbf{V}} = (\omega_j^{k-1})_{j,k \in [n]} \in \mathbb{Z}_q^{n \times n}$.

1.1.5 Modules, Ideals and Units

Let K be a number field with R its ring of integers and $d \in \mathbb{N}$. A subset $\mathcal{M} \subseteq K^d$ is an R -module of rank d if it is closed under addition by elements of \mathcal{M} and under multiplication by elements of R . It is a finitely generated module if there exists a finite family $\{b_k\}_k$ of vectors in K^d such that $\mathcal{M} = \sum_k R \cdot b_k$. An ideal $\mathcal{I} \subseteq R$ is a non-zero additive subgroup of R that is closed under multiplication by R . In other words, ideals are modules of rank $d = 1$. An ideal \mathcal{I} is principal if it is generated by a single element u , meaning $\mathcal{I} = uR = \langle u \rangle$.

An ideal $\mathfrak{p} \neq R$ of R is prime if for all $a, b \in R$, $a \cdot b \in \mathfrak{p}$ implies that a or b is in \mathfrak{p} . We define the norm of an ideal $N(\mathcal{I})$ as the index of \mathcal{I} as an additive subgroup of R , which corresponds to $N(\mathcal{I}) = |R/\mathcal{I}|$.

We define the dual of an ideal \mathcal{I} by $\mathcal{I}^\vee = \{x \in K: \text{Tr}(xy) \in \mathbb{Z} \forall y \in \mathcal{I}\}$ and the dual of a module $\mathcal{M} \subseteq K^d$ by $\mathcal{M}^\vee = \{x \in K^d: \text{Tr}(\langle x, y \rangle) \in \mathbb{Z} \forall y \in \mathcal{M}\}$. Note that R trivially defines an ideal and thus we can define its dual R^\vee . In the case of $R = \mathbb{Z}[\zeta]$, it yields $R^\vee = \frac{1}{f'(\zeta)}R$, where $f(x)$ is the defining polynomial of K . To facilitate notations, we set $\lambda = f'(\zeta)$, such that $\lambda \cdot R^\vee = R$

Example 1.3 (Power-of-2 Cyclotomic Fields)

For the ν -th cyclotomic field K , where ν is a power of two and $n = \nu/2$, it holds $R^\vee = \frac{1}{n}R$. In other words, R^\vee is just a scaling of R .

In the construction of Lemma 2.2 in Chapter 2, we need a condition for elements of R_η for some $\eta \in \mathbb{N}$, i.e., elements of R with a coefficient vector in $\{0, \dots, \eta - 1\}^n$, to be invertible in R_q for a specific prime q . Recall that non-trivial invertible elements in a ring are called units. To do so, we rely on the small norm condition for cyclotomic fields proven by Lyubashevsky and Seiler [LS18].

Lemma 1.2 (Invertibility [LS18, Thm. 1.1])

For $\nu \in \mathbb{N}$, let K be the ν -th cyclotomic field and let $\nu = \prod_j p_j^{e_j}$ be its prime-power factorization, with $e_j \geq 1$. We denote by R the ring of integers of K . Also, let $\mu = \prod_j p_j^{f_j}$ for some $f_j \in [e_j]$. Let q be a prime such that $q \equiv 1 \pmod{\mu}$, and $\text{ord}_\nu(q) = \nu/\mu$, where ord_ν is the multiplicative order modulo ν . Then, any element y of R_q satisfying $0 < \|\tau(y)\|_\infty < q^{1/\varphi(\mu)}/\mathfrak{s}_1(\mu)$ is a unit in R_q , where $\mathfrak{s}_1(\mu)$ denotes the largest singular value of the Vandermonde matrix of the μ -th cyclotomic field.

If ν is a prime power, then so is μ and then [LPR13] states that $\mathfrak{s}_1(\mu) = \sqrt{\mu}$ if μ is odd, and $\mathfrak{s}_1(\mu) = \sqrt{\mu/2}$ otherwise. For more general cases, we refer to the discussions of Lyubashevsky and Seiler [LS18, Conj. 2.6]. We also refer to [LS18, Thm. 2.5] that establishes the density of such primes q for specific values of ν and μ .

📌 Example 1.4 (Power-of-2 Cyclotomic Fields [LS18, Cor. 1.2])

For the ν -th cyclotomic field K , where ν is a power of 2 and $n = \nu/2$, Lemma 1.2 specifies to the following statement. For any power of two $\kappa \leq n$ such that the prime q fulfills $q = 2\kappa + 1 \pmod{4\kappa}$, the invertibility condition for $y \in R_q$ becomes $0 < \|\tau(y)\|_\infty < q^{1/\kappa}/\sqrt{\kappa}$.

We further recall a result from Wang and Wang [WW19] that we need in the proof of Lemma 2.1 in Chapter 2 to construct an invertible matrix over R_q . We say that for $\ell \in [k]$, the vectors $\mathbf{a}_1, \dots, \mathbf{a}_\ell \in R_q^k$ are R_q -linearly independent if for all $x_1, \dots, x_\ell \in R_q$, $\sum_{j \in [\ell]} x_j \mathbf{a}_j = \mathbf{0} \pmod{qR}$ implies $x_1 = \dots = x_\ell = 0$. Further, a matrix $\mathbf{A} \in R_q^{k \times k}$ is invertible modulo qR if there exists a matrix $\mathbf{A}^{-1} \in R_q^{k \times k}$ such that $\mathbf{A} \cdot \mathbf{A}^{-1} = \mathbf{A}^{-1} \cdot \mathbf{A} = \mathbf{I}_k \pmod{qR}$.

Lemma 1.3 (Independence and Invertibility [WW19, Lem. 9+18])

Let K be the ν -th cyclotomic field of degree $n = \varphi(\nu)$, and R its ring of integers. Let $q, k \in \mathbb{N}$ such that q is a prime that verifies $q \geq n$ and $q \nmid \nu$. Then for any $j \in [k-1]$ and R_q -linearly independent vectors $\mathbf{a}_1, \dots, \mathbf{a}_j \in R_q^k$, the probability of sampling a vector $\mathbf{b} \leftarrow U(R_q^k)$ such that $\mathbf{a}_1, \dots, \mathbf{a}_j, \mathbf{b}$ are R_q -linearly independent is at least $1 - \frac{n}{q}$. Let $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_k] \in R_q^{k \times k}$. Then, \mathbf{A} is invertible modulo qR if and only if $\mathbf{a}_1, \dots, \mathbf{a}_k$ are R_q -linearly independent.

1.2 Lattice Theory

In this section, we briefly present the definition of Euclidean and structured lattices together with some of their important properties. For a more detailed introduction of lattices for their use in lattice-based cryptography we refer for instance to [MR09] and [Pei16a].

1.2.1 Euclidean Lattices

An Euclidean lattice Λ is a discrete subgroup of \mathbb{R}^n . Since the space H as introduced in Section 1.1.1 is isomorphic to \mathbb{R}^n , we sometimes consider lattices that are discrete subgroups of H . Each lattice can be represented by a basis matrix $\mathbf{B} = (\mathbf{b}_j)_{j \in [r]} \in \mathbb{R}^{n \times r}$ for some $r \leq n \in \mathbb{N}$, as the set of all integer linear combinations of the basis elements, i.e., $\Lambda = \sum_{j \in [r]} \mathbb{Z} \cdot \mathbf{b}_j$. The dimension of the lattice is n and the rank is r . In this thesis, we only consider full-rank lattices, namely lattices for which $r = n$.

The origin-centered fundamental parallelepiped $\mathcal{P}(\mathbf{B})$ of the lattice Λ generated by the basis $\mathbf{B} = (\mathbf{b}_j)_{j \in [n]}$ is defined by $\mathcal{P}(\mathbf{B}) = \left\{ \sum_{j \in [n]} z_j \mathbf{b}_j : z_j \in \left[-\frac{1}{2}, \frac{1}{2}\right), \forall j \in [n] \right\}$. For any $\mathbf{w} \in \mathbb{R}^n$, we write $\mathbf{x} = \mathbf{w} \pmod{\mathbf{B}}$ or $\mathbf{x} = \mathbf{w} + \Lambda$ to denote the unique point $\mathbf{x} \in \mathcal{P}(\mathbf{B})$ such that $\mathbf{w} - \mathbf{x} \in \Lambda$. The volume of $\mathcal{P}(\mathbf{B})$ is called the determinant of Λ and denoted by $\det(\Lambda)$. The determinant is independent of the choice of the basis \mathbf{B} and thus an invariant of the lattice. For any basis \mathbf{B} of a full-rank lattice Λ it yields $\det(\Lambda) = |\det(\mathbf{B})|$.

We define the dual lattice of a lattice Λ by $\Lambda^* = \{\mathbf{x} \in \text{Span}(\Lambda) : \forall \mathbf{y} \in \Lambda, \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$. If \mathbf{B} is a basis of Λ , then $\mathbf{B}^* = (\mathbf{B}^T)^{-1}$ is a basis of Λ^* . The first minimum $\lambda_1(\Lambda)$ of a lattice Λ is the Euclidean norm of any of its shortest non-zero vectors. The first minimum with respect to the infinity norm is denoted by $\lambda_1^\infty(\Lambda)$. We further define the second minimum $\lambda_2(\Lambda)$ as the norm of any second shortest non-zero vector, which we require to be linearly independent of all shortest vectors.

Using Minkowski's theorem, we can upper bound the norm of a shortest non-zero vector in arbitrary lattices. More concretely, for a lattice Λ of dimension n and determinant $\det(\Lambda)$, it yields $\lambda_1(\Lambda) \leq \sqrt{n} \det(\Lambda)^{1/n}$. Heuristically, one can estimate the expected norm of a shortest non-zero vector in randomly chosen lattices by using the Gaussian heuristic, which slightly improves the Minkowski bound. It says that for an n -dimensional lattice Λ with determinant $\det(\Lambda)$, we expect

$$\lambda_1(\Lambda) \approx \sqrt{\frac{n}{2\pi e}} \cdot \det(\Lambda)^{1/n}. \quad (1.2)$$

There is a special class of lattices which plays an important role in lattice-based cryptography. Given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ for some integers $n, m, q \in \mathbb{N}$, we can define two lattices

$$\begin{aligned} \Lambda_q(\mathbf{A}) &= \{\mathbf{y} \in \mathbb{Z}^m : \mathbf{y} = \mathbf{A}^T \mathbf{s} \bmod q \text{ for some } \mathbf{s} \in \mathbb{Z}^n\} \quad \text{and} \\ \Lambda_q^\perp(\mathbf{A}) &= \{\mathbf{y} \in \mathbb{Z}^m : \mathbf{A}\mathbf{y} = \mathbf{0} \bmod q\}. \end{aligned}$$

Both lattices are of dimension m . The first is generated by the rows of \mathbf{A} , whereas the second contains all vectors that are orthogonal to the rows of \mathbf{A} . Furthermore, they are connected via lattice duality, i.e., $\Lambda_q^\perp(\mathbf{A}) = q \cdot \Lambda_q(\mathbf{A})^*$ and $\Lambda_q(\mathbf{A}) = q \cdot \Lambda_q^\perp(\mathbf{A})^*$.

1.2.2 Structured Lattices

We now define two types of structured lattices. Let K be number field with ring of integers R .

To define the first class of structured lattices, we observe that any ideal \mathcal{I} over R embeds into a lattice $\sigma(\mathcal{I})$ in H , and a lattice $\sigma_H(\mathcal{I})$ in \mathbb{R}^n , which we call *ideal lattices*. The ideal lattice inherits the algebraic structure of the corresponding ideal, i.e., multiplying a lattice vector by any ring element maps it again to a lattice vector. This is not true for general Euclidean lattice.

Regarding the second class of structured lattices, we remark that for an R -module $\mathcal{M} \subseteq K^d$, the set $(\sigma, \dots, \sigma)(\mathcal{M})$ is a lattice in H^d and $(\sigma_H, \dots, \sigma_H)(\mathcal{M})$ is a lattice in \mathbb{R}^{nd} , both of which are called *module lattices*. The positive integer d is the module rank and module lattices of rank 1 are in fact ideal lattices. Again, the module lattice inherits the algebraic structure of the underlying module, making it closed with respect to scalar multiplication by ring elements.

To ease readability, we simply use \mathcal{I} (resp. \mathcal{M}) to denote the ideal lattice (resp. the module lattice). Note that the ideal lattice $\sigma(\mathcal{I}^\vee)$ corresponding to the ideal \mathcal{I} and the module lattice $\sigma(\mathcal{M}^\vee)$ corresponding to the module \mathcal{M} are the same as the respective dual lattices up to complex conjugation, i.e., $\sigma(\mathcal{I}^\vee) = \overline{\sigma(\mathcal{I})}^*$ and $\sigma(\mathcal{M}^\vee) = \overline{\sigma(\mathcal{M})}^*$.

Recall that for arbitrary lattices Minkowski's theorem provides an upper bound on their first minimum. For ideal (but not module) lattices, we can use the geometric-arithmetic mean inequality to additionally give a lower bound for their minimum. Here, we state the result with respect to the infinity norm, but it can be shown for any ℓ_p -norm.

Lemma 1.4 (Lower Bound on Minima [PR07, Lem. 6.2])

For any ideal \mathcal{I} over a number field K with ring of integers R of degree n , it yields

$$\lambda_1^\infty(\sigma_H(\mathcal{I})) \geq N(\mathcal{I})^{1/n}.$$

We also note that for the infinity norm, the first minimum of the module lattice $\mathcal{I}^d = \mathcal{I} \times \dots \times \mathcal{I}$ is the same as the first minimum of the ideal lattice \mathcal{I} , i.e., $\lambda_1^\infty(\mathcal{I}^d) = \lambda_1^\infty(\mathcal{I})$.

1.2.3 Computational Problems

One of the most studied lattice problems is the Shortest Vector Problem (SVP). It exists in both search and decision versions, which we define in their approximate variants in the following.

Definition 1.1 (Search Shortest Vector Problem): Let $\gamma = \gamma(n) \geq 1$ be a function in the dimension n . An input to the search Shortest Vector Problem SVP_γ is a basis \mathbf{B} of an n -dimensional lattice Λ . The goal is to find a vector $\mathbf{z} \neq \mathbf{0}$ such that $\|\mathbf{z}\|_2 \leq \gamma \cdot \lambda_1(\Lambda)$.

The SVP restricted to ideal lattice is denoted by Id-SVP_γ and to module lattices by Mod-SVP_γ .

Definition 1.2 (Decision Shortest Vector Problem): Let $\gamma = \gamma(n) \geq 1$ be a function in the dimension n . An input to the decision Shortest Vector Problem GapSVP_γ is a pair (\mathbf{B}, δ) , where \mathbf{B} is a basis of an n -dimensional lattice Λ and $\delta > 0$ is a real number. It is a YES instance if $\lambda_1(\Lambda) \leq \delta$, and it is a NO instance if $\lambda_1(\Lambda) > \gamma \cdot \delta$. The problem asks to distinguish whether a given instance is a YES or a NO instance. If $\lambda_1(\Lambda)$ falls in the interval $(\delta, \gamma \cdot \delta]$, any answer is correct.

The GapSVP_γ restricted to module lattices is denoted by Mod-GapSVP_γ . Whereas the problem GapSVP_γ is conjectured to be hard to solve for γ polynomial in the lattice dimension, the problem Mod-GapSVP_γ becomes easy in the special case of modules of rank 1 (i.e., ideals) as the minimum of ideal lattices can be bounded above *and* below, see Lemma 1.4. However, it is still conjectured that Mod-GapSVP_γ is hard to solve for γ polynomial in the module lattice dimension if the module has rank at least 2.

Another lattice problem is the Unique Shortest Vector Problem (u-SVP), where we need to find a shortest vector, having the additional promise that there is a gap between the first and second minimum of the lattice.

Definition 1.3 (Unique Shortest Vector Problem): Let $\delta \geq 1$. An input to the Unique Shortest Vector Problem u-SVP_δ is given by a basis \mathbf{B} of an n -dimensional lattice Λ such that $\lambda_2(\Lambda) \geq \delta \cdot \lambda_1(\Lambda)$. The goal is to find a vector $\mathbf{x} \in \Lambda$ of norm $\lambda_1(\Lambda)$.

If $\delta = 1$, this problem becomes the exact SVP problem from above (i.e., Def. 1.1 with $\gamma = 1$). Note that the problem becomes easier for increasing δ .

Within this thesis, we further need two intermediate lattice problems, presented in the following. The first comes in two flavors, one for general Euclidean lattices, and one for module lattices. For technical reasons, we state the first one with respect to the infinity norm, whereas in the second one we use the $(2, \infty)$ -norm as defined in Section 1.1.2.

Definition 1.4 (Bounded Distance Decoding): Let \mathbf{B} be a basis of an n -dimensional lattice $\Lambda(\mathbf{B})$ and δ be a positive real. An input to the Bounded Distance Decoding problem BDD_δ is a point $\mathbf{y} \in \mathbb{R}^n$ of the form $\mathbf{y} = \mathbf{x} + \mathbf{e}$, where $\mathbf{x} \in \Lambda(\mathbf{B})$ and $\|\mathbf{e}\|_\infty \leq \delta$. The problem asks to find \mathbf{x} (or equivalently \mathbf{e}).

If $\delta < \lambda_1(\Lambda(\mathbf{B}))/2$, then the solution to BDD_δ is unique.

Definition 1.5 (Module Bounded Distance Decoding): Let K be a number field with R its ring of integers of degree n and $\mathcal{M} \subseteq K^d$ be a module of R of rank d . Further, let δ be a positive real number. An input to the Module Bounded Distance Decoding problem $\text{BDD}_{\mathcal{M},\delta}$ is a point $\mathbf{y} \in K^d$ of the form $\mathbf{y} = \mathbf{x} + \mathbf{e}$, where $\mathbf{x} \in \mathcal{M}$ and $\|\mathbf{e}\|_{2,\infty} \leq \delta$. The problem asks to find \mathbf{x} (or equivalently \mathbf{e}).

By D_g we denote the continuous Gaussian distribution of width g over $K_{\mathbb{R}}^d$, which we define properly in Section 1.3.2. As we only need the following problem for module lattices, we directly define it in this setting.

Definition 1.6 (Gaussian Decoding Problem): Let K be a number field with R its ring of integers of degree n and $\mathcal{M} \subseteq K^d$ be a module of R of rank d . Further, let $g > 0$ be a Gaussian parameter. An input to the Gaussian Decoding Problem $\text{GDP}_{\mathcal{M},g}$ is a coset $\mathbf{e} + \mathcal{M}$, where $\mathbf{e} \leftarrow D_g$. The goal is to find \mathbf{e} .

Every $\text{GDP}_{\mathcal{M},g}$ instance defines a $\text{BDD}_{\mathcal{M},\delta}$ instance with $\delta = g \cdot \sqrt{d} \cdot \omega(\sqrt{\log_2 n})$. Note that all computational problems can also be defined with regard to other norms. We finally recall (for the module setting) the well-known result that the LLL algorithm together with Babai's round-off algorithm solves BDD for an exponentially large δ .

Lemma 1.5 (The LLL Algorithm [LLL82, Bab85])

Let K be a number field with R its ring of integers of degree n and $\mathcal{M} \subseteq K^d$ be a module of R of rank d . There exists a polynomial-time algorithm that finds a vector of norm at most $2^{\frac{N}{2}} \cdot \lambda_1(\mathcal{M})$ and solves $\text{BDD}_{\mathcal{M},\delta}$ for $\delta = 2^{-\frac{N}{2}} \cdot \lambda_1(\mathcal{M})$, where $N = nd$.

1.3 Probability Theory

In this section, we introduce the notions of probability theory needed within this manuscript. First, we recall the definitions and some properties of Gaussian measures and how to define them over number fields. Then, we define two measures of distance for probability distributions that are commonly used in cryptography. We conclude this section by providing two flavors of the so-called Leftover Hash Lemma.

For a finite set S , we denote its cardinality by $|S|$ and the uniform distribution over S by $U(S)$. The operation of sampling an element $x \in S$ according to a distribution D over S is denoted by $x \leftarrow D$, where the set S is implicit.

1.3.1 Gaussian Measures

For a positive definite matrix $\mathbf{\Sigma} \in \mathbb{R}^{n \times n}$, a vector $\mathbf{c} \in \mathbb{R}^n$, we define the Gaussian function by $\rho_{\mathbf{c},\sqrt{\mathbf{\Sigma}}}(\mathbf{x}) = \exp(-\pi(\mathbf{x} - \mathbf{c})^T \mathbf{\Sigma}^{-1}(\mathbf{x} - \mathbf{c}))$ for all $\mathbf{x} \in \mathbb{R}^n$. We extend this definition to the degenerate case, i.e., positive semi-definite, by considering the generalized Moore-Penrose inverse. For convenience, we use the same notation as the standard inverse. We then define the continuous Gaussian probability distribution by its density $D_{\mathbf{c},\sqrt{\mathbf{\Sigma}}}(\mathbf{x}) = (\det(\mathbf{\Sigma}))^{-1/2} \rho_{\mathbf{c},\sqrt{\mathbf{\Sigma}}}(\mathbf{x})$. By abuse of notation, we call $\mathbf{\Sigma}$ the covariance matrix, even if in theory the covariance matrix of $D_{\mathbf{c},\sqrt{\mathbf{\Sigma}}}$ is $\mathbf{\Sigma}/(2\pi)$. If $\mathbf{\Sigma}$ is diag-




Figure 1.1: Graph of the probability density function of a continuous Gaussian in dimension $n = 1$ and of the probability mass function of a discrete Gaussian over the lattice $\Lambda = \mathbb{Z}$.

onal with diagonal vector $\mathbf{r}^2 \in (\mathbb{R}^+)^n$, we simply write $D_{\mathbf{c}, \mathbf{r}}$, and if $\mathbf{c} = 0$, we omit it. When $\Sigma = s^2 \mathbf{I}_n$ for some positive real s , we simplify further to $D_{\mathbf{c}, s}$. We then define the discrete Gaussian distribution over a lattice Λ by conditioning \mathbf{x} to be in Λ , i.e., $\mathcal{D}_{\Lambda, \mathbf{c}, \sqrt{\Sigma}}(\mathbf{x}) = D_{\mathbf{c}, \sqrt{\Sigma}}(\mathbf{x}) / D_{\mathbf{c}, \sqrt{\Sigma}}(\Lambda)$ for all $\mathbf{x} \in \Lambda$, where $D_{\mathbf{c}, \sqrt{\Sigma}}(\Lambda) = \sum_{\mathbf{y} \in \Lambda} D_{\mathbf{c}, \sqrt{\Sigma}}(\mathbf{y})$. For the simple one-dimensional case, we sketch their corresponding probability and mass functions in Figure 1.1. An important property of continuous and discrete Gaussians is that they allow for concrete tail bounds. In this thesis, we only need the tail bound for the discrete case, as stated in the following.

Lemma 1.6 (Discrete Gaussian Tail Bound [Ban93, Lem. 1.5])

For any lattice $\Lambda \subseteq \mathbb{R}^n$, vector $\mathbf{c} \in \mathbb{R}^n$ and parameter $s > 0$, it yields

$$\Pr_{\mathbf{x} \leftarrow \mathcal{D}_{\Lambda, \mathbf{c}, s}} [\|\mathbf{x} - \mathbf{c}\|_2 \leq s\sqrt{n}] \geq 1 - 2^{-\Omega(n)}.$$

The smoothing parameter of a lattice Λ , denoted by $\eta_\varepsilon(\Lambda)$ for some $\varepsilon > 0$ and introduced by Micciancio and Regev [MR07], is the smallest $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$. It represents the smallest Gaussian parameter $s > 0$ such that the discrete Gaussian $\mathcal{D}_{\Lambda, \mathbf{c}, s}$ behaves in many respects like a continuous Gaussian distribution. We recall the following bounds and properties of the smoothing parameter that we need throughout this thesis.

Lemma 1.7 ([Pei08, Lem. 3.5])

For an n -dimensional lattice Λ and $\varepsilon > 0$, we have

$$\eta_\varepsilon(\Lambda) \leq \sqrt{\ln(2n(1 + 1/\varepsilon)) / \pi} / \lambda_1^\infty(\Lambda^*).$$

Whereas the first bound uses the minimum with respect to the infinity norm, the next uses the minimum with respect to the Euclidean norm.

Lemma 1.8 ([Ban93, Lem. 1.5] and [Reg05, Claim 2.13])

Let Λ be an n -dimensional lattice and $\varepsilon = \exp(-n)$, it holds

$$\frac{\sqrt{n}}{\sqrt{\pi} \lambda_1(\Lambda^*)} \leq \eta_\varepsilon(\Lambda) \leq \frac{\sqrt{n}}{\lambda_1(\Lambda^*)}.$$

Lemma 1.9 ([MR07, Lem. 4.1])

Let Λ be an n -dimensional lattice, $\varepsilon > 0$, and $s > \eta_\varepsilon(\Lambda)$. Then the distribution of the coset $\mathbf{e} + \Lambda$, where $\mathbf{e} \leftarrow D_s$, is within statistical distance $\varepsilon/2$ of the uniform distribution over the cosets of Λ .

Recall that continuous Gaussian distributions are linear and in particular the sum of two Gaussians $D_{\mathbf{c}_1, \sqrt{\Sigma_1}}$ and $D_{\mathbf{c}_2, \sqrt{\Sigma_2}}$ follows the Gaussian distribution $D_{\mathbf{c}_1 + \mathbf{c}_2, \sqrt{\Sigma}}$, where $\Sigma = \Sigma_1 + \Sigma_2$. To obtain similar results for the discrete case, we need the smoothing parameter. In the following, we consider the special case of the sum of a continuous Gaussian and a discrete one. In particular, the lemma generalizes known results to elliptical Gaussians, which we use in the proof of Lemma 1.16.

Lemma 1.10 (Adapted from [LS15, Lem. 2.8] & [Reg09, Claim 3.9])

Let Λ be an n -dimensional lattice, $\mathbf{a} \in \mathbb{R}^n$, \mathbf{R}, \mathbf{S} two positive semi-definite matrices of $\mathbb{R}^{n \times n}$, and $\mathbf{T} = \mathbf{R} + \mathbf{S}$. We also define $\mathbf{U} = (\mathbf{R}^{-1} + \mathbf{S}^{-1})^{-1}$, and we assume that $\rho_{\sqrt{\mathbf{U}^{-1}}}(\Lambda^* \setminus \{0\}) \leq \varepsilon$ for some $\varepsilon \in (0, 1/2)$. Consider the distribution Y over \mathbb{R}^n obtained by adding a discrete sample from $\mathcal{D}_{\Lambda + \mathbf{a}, \sqrt{\mathbf{R}}}$ and a continuous sample from $D_{\sqrt{\mathbf{S}}}$. Then we have $\Delta(Y, D_{\sqrt{\mathbf{T}}}) \leq 2\varepsilon$.

Proof: The density function of the distribution Y is given by

$$\begin{aligned}
Y(\mathbf{x}) &= \sum_{\mathbf{y} \in \Lambda + \mathbf{a}} \mathcal{D}_{\Lambda + \mathbf{a}, \sqrt{\mathbf{R}}}(\mathbf{y}) D_{\sqrt{\mathbf{S}}}(\mathbf{x} - \mathbf{y}) \\
&= \frac{1}{\rho_{-\mathbf{a}, \sqrt{\mathbf{R}}}(\Lambda) \sqrt{\det \mathbf{S}}} \sum_{\mathbf{y} \in \Lambda + \mathbf{a}} \rho_{\sqrt{\mathbf{R}}}(\mathbf{y}) \rho_{\sqrt{\mathbf{S}}}(\mathbf{x} - \mathbf{y}) \\
&= \frac{1}{\rho_{-\mathbf{a}, \sqrt{\mathbf{R}}}(\Lambda) \sqrt{\det \mathbf{S}}} \sum_{\mathbf{y} \in \Lambda + \mathbf{a}} \rho_{\sqrt{\mathbf{T}}}(\mathbf{x}) \rho_{\mathbf{R}\mathbf{T}^{-1}\mathbf{x}, \sqrt{\mathbf{U}}}(\mathbf{y}) \quad [\text{Pei10, Fact 2.1}] \\
&= \frac{\rho_{\sqrt{\mathbf{T}}}(\mathbf{x})}{\sqrt{\det \mathbf{T}}} \cdot \frac{\sqrt{\det \mathbf{T}} \rho_{\mathbf{R}\mathbf{T}^{-1}\mathbf{x}, \sqrt{\mathbf{U}}}(\Lambda)}{\sqrt{\det \mathbf{T}} \rho_{-\mathbf{a}, \sqrt{\mathbf{R}}}(\Lambda)} \\
&= D_{\sqrt{\mathbf{T}}}(\mathbf{x}) \cdot \frac{(\sqrt{\det \mathbf{R}} \sqrt{\det \mathbf{S}} / \sqrt{\det \mathbf{T}})^{-1} \widehat{\rho_{\mathbf{x}', \sqrt{\mathbf{U}}}(\Lambda^*)}}{(\sqrt{\det \mathbf{R}})^{-1} \widehat{\rho_{-\mathbf{a}, \sqrt{\mathbf{R}}}(\Lambda^*)}},
\end{aligned}$$

where $\mathbf{x}' = \mathbf{R}\mathbf{T}^{-1}\mathbf{x}$, and \widehat{f} denotes the Fourier transform of f . First notice that $(\det \mathbf{R} \cdot \det \mathbf{S}) / \det \mathbf{T} = 1 / \det(\mathbf{R}^{-1}\mathbf{T}\mathbf{S}^{-1}) = 1 / \det \mathbf{U}^{-1}$. Moreover, recalling that $\widehat{\rho_{\mathbf{c}, \sqrt{\Sigma}}(\mathbf{w})} = \sqrt{\det \Sigma} e^{-2i\pi \langle \mathbf{c}, \mathbf{w} \rangle} \rho_{\sqrt{\Sigma^{-1}}}(\mathbf{w})$, we get

$$\left| 1 - (\sqrt{\det \mathbf{U}})^{-1} \widehat{\rho_{\mathbf{x}', \sqrt{\mathbf{U}}}(\Lambda^*)} \right| \leq \rho_{\sqrt{\mathbf{U}^{-1}}}(\Lambda^* \setminus \{0\}) \leq \varepsilon.$$

For the denominator, we first notice that for two positive semi-definite matrices \mathbf{A} and \mathbf{B} , if $\mathbf{A} - \mathbf{B}$ is positive semi-definite, then $\rho_{\sqrt{\mathbf{A}}}(\mathbf{w}) \geq \rho_{\sqrt{\mathbf{B}}}(\mathbf{w})$ for all $\mathbf{w} \in \mathbb{R}^n$. Since $\mathbf{U}^{-1} - \mathbf{R}^{-1} = \mathbf{S}^{-1}$ is positive semi-definite, it yields $\rho_{\sqrt{\mathbf{R}^{-1}}}(\Lambda^* \setminus \{0\}) \leq \rho_{\sqrt{\mathbf{U}^{-1}}}(\Lambda^* \setminus \{0\}) \leq \varepsilon$. Therefore, using

the same method as above, we get

$$\left| 1 - (\sqrt{\det \mathbf{R}})^{-1} \widehat{\rho_{-\mathbf{a}, \sqrt{\mathbf{R}}}(\Lambda^*)} \right| \leq \rho_{\sqrt{\mathbf{R}^{-1}}}(\Lambda^* \setminus \{0\}) \leq \varepsilon.$$

which leads to

$$\frac{(\sqrt{\det \mathbf{R}} \sqrt{\det \mathbf{S}} / \sqrt{\det \mathbf{T}})^{-1} \widehat{\rho_{\mathbf{x}', \sqrt{\mathbf{U}}}(\Lambda^*)}}{(\sqrt{\det \mathbf{R}})^{-1} \widehat{\rho_{-\mathbf{a}, \sqrt{\mathbf{R}}}(\Lambda^*)}} \in \left[\frac{1 - \varepsilon}{1 + \varepsilon}, \frac{1 + \varepsilon}{1 - \varepsilon} \right] \subseteq [1 - 2\varepsilon, 1 + 4\varepsilon],$$

assuming that $\varepsilon < 1/2$. We thus end up with $|Y(\mathbf{x}) - D_{\sqrt{\mathbf{T}}}(\mathbf{x})| \leq 4\varepsilon D_{\sqrt{\mathbf{T}}}(\mathbf{x})$. Integration and factor 1/2 of the statistical distance yield the lemma. \blacksquare

Lemma 1.11 ([BLP⁺13, Lem. 2.10] & [Pei10, Thm. 3.1])

Let Λ be an n -dimensional lattice, $\varepsilon \in (0, 1/2]$, and $\beta, r > 0$ such that $r \geq \eta_\varepsilon(\Lambda)$. Then the distribution of $\mathbf{x} + \mathbf{y}$, obtained by first sampling \mathbf{x} from D_β , and then sampling \mathbf{y} from $\mathcal{D}_{\Lambda, \mathbf{x}, r}$, is within statistical distance 8ε of $\mathcal{D}_{\Lambda, \sqrt{\beta^2 + r^2}}$.

Note that in Lemma 1.11 the parameter of the discrete Gaussian distribution of \mathbf{y} depends on the vector \mathbf{x} , sampled from a continuous Gaussian beforehand. This is not the case in Lemma 1.10.

Finally, in Chapter 3 we need an exact sampler of discrete Gaussians over lattices. For a basis $\mathbf{B} = (\mathbf{b}_j)_{j \in [n]}$ of an n -dimensional lattice, we denote by $\text{GS}(\mathbf{B}) = (\text{GS}(\mathbf{b}_j))_{j \in [n]}$ the Gram-Schmidt orthogonalization of \mathbf{B} from left to right.

Lemma 1.12 (Gaussian Sampler [GPV08, Thm. 4.1] & [BLP⁺13, Lem. 2.3])

Let $n \in \mathbb{N}$. There exists a PPT algorithm \mathcal{D} that, given a basis $\mathbf{B} = (\mathbf{b}_j)_{j \in [n]}$ of an n -dimensional lattice Λ , a parameter $r \geq \max_{j \in [n]} \|\text{GS}(\mathbf{b}_j)\|_2 \cdot \sqrt{\ln(2n+4)/\pi}$ and a center $\mathbf{c} \in \mathbb{R}^n$, outputs a sample whose distribution is $D_{\Lambda, r, \mathbf{c}}$.

1.3.2 Gaussians over Number Fields

In this section, we introduce Gaussian distributions over $K_{\mathbb{R}}^d$, where $K = \mathbb{Q}(\zeta)$ is a number field, R its ring of integers, and $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$.² We further prove some of their properties needed in the rest of this thesis. Gaussian distributions over $K_{\mathbb{R}}$ have been introduced alongside the R-LWE problem [LPR10], and then generalized and used in most papers dealing with structured variants of LWE.³ Let n denote the degree of the number field K and let $d \in \mathbb{N}$. Further, recall the definition of the canonical embedding σ_H as introduced in Section 1.1.2. We define general Gaussian distributions over $K_{\mathbb{R}}^d$ through their canonical embedding σ_H to \mathbb{R}^{nd} , namely sampling $\mathbf{y}^{(H)} \in \mathbb{R}^{nd}$ according to $D_{\sqrt{\mathbf{\Sigma}}}$ for some positive semi-definite matrix $\mathbf{\Sigma}$ in $\mathbb{R}^{nd \times nd}$ and then mapping it back to $K_{\mathbb{R}}^d$ by $\mathbf{y} = \sigma_H^{-1}(\mathbf{y}^{(H)})$. To ease readability, we denote the described distribution of $\mathbf{y} \in K_{\mathbb{R}}^d$ by $D_{\sqrt{\mathbf{\Sigma}}}$. As for the Gaussians over \mathbb{R} , we write $D_{\mathbf{r}}$ if $\mathbf{\Sigma}$ is a diagonal matrix with diagonal

²We need to consider the real tensor field $K_{\mathbb{R}}$ as the canonical embedding is an isomorphism between $K_{\mathbb{R}}$ and H but not between R and H , nor K and H .

³We refer to Section 1.4 for a proper introduction of LWE and its structured variants.

vector $\mathbf{r}^2 \in (\mathbb{R}^+)^{nd}$. For $0 \leq \alpha < \alpha'$, we define $\Psi_{[\alpha, \alpha']}$ to be the set of Gaussian distributions $D_{\mathbf{r}}$ with $\alpha < \|\mathbf{r}\|_{\infty} \leq \alpha'$. If $\alpha = 0$, we simply write $\Psi_{\leq \alpha'}$.

Sometimes, it is more convenient to consider a Gaussian distribution $D_{\sqrt{\Sigma}}$ over $K_{\mathbb{R}}$ and extend it for $d \in \mathbb{N}$ to a Gaussian distribution over $K_{\mathbb{R}}^d$ by simply drawing all d coefficients independently over $K_{\mathbb{R}}$. In this case, we write $(D_{\sqrt{\Sigma}})^d$.

Further, for any positive real α , we define the specific distribution Υ_{α} over distributions on H as done by Peikert et al. [PRS17a]. Let $n = t_1 + 2t_2$ be the decomposition in real and complex embeddings and fix an arbitrary function $g(n) = \omega(\sqrt{\log_2 n})$. A distribution sampled from Υ_{α} is an elliptical Gaussian $D_{\mathbf{r}}$, where \mathbf{r} is sampled as follows: For $j \in [t_1]$, sample $x_j \leftarrow D_1$ and set $r_j^2 = \alpha^2(x_j^2 + g^2(n))/2$. For $j \in \{t_1 + 1, \dots, t_1 + t_2\}$, sample $x_j, y_j \leftarrow D_{1/\sqrt{2}}$ and set $r_j^2 = r_{j+t_2}^2 = \alpha^2(x_j^2 + y_j^2 + g^2(n))/2$. We need this distribution for the classical reduction in Chapter 3.

In the proof of Lemma 2.3, we need to identify the distribution of $\mathbf{y} = \mathbf{U}\mathbf{e}$ for an arbitrary matrix \mathbf{U} and a Gaussian vector $\mathbf{e} \in K_{\mathbb{R}}^d$ for which the components are independent of each other. To do so, we define for $m, d \in \mathbb{N}$ the ring homomorphism $\theta : K_{\mathbb{R}}^{m \times d} \rightarrow \mathbb{C}^{nm \times nd}$ given by

$$\theta(\mathbf{A}) = \begin{bmatrix} \mathbf{D}_{1,1} & \cdots & \mathbf{D}_{1,d} \\ \vdots & \ddots & \vdots \\ \mathbf{D}_{m,1} & \cdots & \mathbf{D}_{m,d} \end{bmatrix},$$

with $\mathbf{A} = (a_{jk})_{j \in [m], k \in [d]}$ and $\mathbf{D}_{j,k} = \text{diag}(\sigma(a_{jk})) \in \mathbb{C}^{n \times n}$.

Lemma 1.13

Let $d, n \in \mathbb{N}$, such that K is a number field of degree n . Let $\mathbb{S} \in \mathbb{R}^{nd \times nd}$ be a positive semi-definite matrix, and $\mathbf{U} \in K_{\mathbb{R}}^{d \times d}$. We define $\Sigma = \left(\mathbb{H}^{\dagger} \theta(\mathbf{U}) \mathbb{H} \right) \mathbb{S} \left(\mathbb{H}^{\dagger} \theta(\mathbf{U}) \mathbb{H} \right)^{\dagger} \in \mathbb{R}^{nd \times nd}$, where $\mathbb{H} = \text{diag}(\mathbf{H}, \dots, \mathbf{H}) \in \mathbb{C}^{nd \times nd}$, with \mathbf{H} the matrix form of the basis of the space H , defined in Section 1.1.1. Then, the distribution of $\mathbf{y} = \mathbf{U}\mathbf{e}$, where $\mathbf{e} \in K_{\mathbb{R}}^d$ is distributed according to $D_{\sqrt{\mathbb{S}}}$, is exactly $D_{\sqrt{\Sigma}}$.

Proof: Let $\mathbf{e} = (e_j)_{j \in [d]} \in K_{\mathbb{R}}^d$ be a Gaussian vector distributed according to $D_{\sqrt{\mathbb{S}}}$. Further, let $\mathbf{U} = (u_{jk})_{j,k \in [d]}$. Our goal is to characterize the distribution of the vector $\mathbf{y} = (y_j)_{j \in [d]}$, which is given as the product of \mathbf{U} and \mathbf{e} . For all $j \in [d]$, we have $y_j = \sum_{k \in [d]} u_{jk} e_k$ and thus $\sigma(y_j) = \sum_{k \in [d]} \sigma(u_{jk}) \circ \sigma(e_k)$, where \circ denotes the component-wise product. The component-wise product $\mathbf{a} \circ \mathbf{b}$ of two vectors \mathbf{a} and \mathbf{b} can also be expressed as the matrix-vector product $\text{diag}(\mathbf{a}) \cdot \mathbf{b}$. It results in

$$\sigma(\mathbf{y}) = \begin{bmatrix} \sigma(y_1) \\ \vdots \\ \sigma(y_d) \end{bmatrix} = \theta(\mathbf{U})\sigma(\mathbf{e}),$$

where $\theta(\mathbf{U})$ is the block matrix $[\text{diag}(\sigma(u_{jk}))]_{j,k \in [d]} \in \mathbb{C}^{nd \times nd}$. As we have seen before, we can decompose σ in the basis of H and get $\sigma(y_j) = \mathbf{H}\mathbf{y}_j^{(H)}$ (respectively $\sigma(e_j) = \mathbf{H}\mathbf{e}_j^{(H)}$) for all $j \in [d]$.

By using the block matrix product, we end up with

$$\sigma(\mathbf{y}) = \begin{bmatrix} \mathbf{H} & & \\ & \ddots & \\ & & \mathbf{H} \end{bmatrix} \begin{bmatrix} \mathbf{y}_1^{(H)} \\ \vdots \\ \mathbf{y}_d^{(H)} \end{bmatrix} = \mathbb{H}\mathbf{y}^{(H)}.$$

Thus $\mathbb{H}\mathbf{y}^{(H)} = \theta(\mathbf{U})\mathbb{H}\mathbf{e}^{(H)}$, which leads to $\mathbf{y}^{(H)} = \mathbb{H}^\dagger\theta(\mathbf{U})\mathbb{H}\mathbf{e}^{(H)}$. Now notice that the blocks of $\mathbb{H}^\dagger\theta(\mathbf{U})\mathbb{H}$ are the $\mathbf{H}^\dagger\text{diag}(\sigma(u_{jk}))\mathbf{H}$ which correspond to the matrix form of the multiplication by u_{jk} in the basis of the space H and thus is in $\mathbb{R}^{n \times n}$. Hence $\mathbb{H}^\dagger\theta(\mathbf{U})\mathbb{H} \in \mathbb{R}^{nd \times nd}$.

By definition, $\mathbf{e}^{(H)}$ is distributed according to $D_{\sqrt{\mathbb{S}}}$. Thus $\mathbf{y}^{(H)}$ is also distributed along a 0-centered Gaussian over \mathbb{R}^{nd} , but with covariance matrix $\boldsymbol{\Sigma} = \left(\mathbb{H}^\dagger\theta(\mathbf{U})\mathbb{H}\right) \mathbb{S} \left(\mathbb{H}^\dagger\theta(\mathbf{U})\mathbb{H}\right)^\dagger$. ■

In particular, when the matrix \mathbb{S} is of the form $\mathbb{S} = \text{diag}(r_1^2, \dots, r_1^2, \dots, r_d^2, \dots, r_d^2)$ for some positive reals r_1, \dots, r_d , then $\sqrt{\mathbb{S}}$ commutes with \mathbb{H} and the covariance simplifies to $\boldsymbol{\Sigma} = \mathbb{H}^\dagger\tilde{\mathbf{U}}\tilde{\mathbf{U}}^\dagger\mathbb{H}$, with $\tilde{\mathbf{U}} = (\text{diag}(\sigma(r_k u_{jk}))_{j,k \in [d]})$.

To apply the noise flooding argument in the proof of the hardness of M-LWE using a binary secret in Section 2.2, we need the following bound on the norm of the product of some discrete Gaussian matrix (in the canonical embedding σ) and a vector of small norm in the dual ring (in the coefficient embedding τ). The result is restricted to cyclotomic number fields, i.e., $K = \mathbb{Q}(\zeta)$, where ζ is a primitive root of unity. As stated in Section 1.1.5, for every number field $K = \mathbb{Q}(\zeta)$ with ring of integers $R = \mathbb{Z}[\zeta]$, there exists a λ such that $\lambda \cdot R^\vee = R$. This implies $\lambda \cdot (R_\eta^\vee)^d = (R_\eta)^d$ for every $\eta, d \in \mathbb{N}$.

Lemma 1.14 (Noise Flooding Bound)

Let $K = \mathbb{Q}(\zeta)$ be a cyclotomic field with $R = \mathbb{Z}[\zeta]$ its ring of integers of degree n and its dual R^\vee . Let $d, m, q, \eta \in \mathbb{N}$ and $\alpha \in (0, 1)$. Sample $\mathbf{Z} \leftarrow (\mathcal{D}_{R^\vee, \alpha q})^{m \times d}$ and $\mathbf{s} \leftarrow U((R_\eta^\vee)^d)$. We set $\tilde{\mathbf{s}} = \lambda \mathbf{s} \in R_\eta^d$. Then, with overwhelming probability $\|\mathbf{Z}\tilde{\mathbf{s}}\|_2 \leq \alpha q n^2 d(\eta-1)\sqrt{m}$. In particular, the Euclidean norm of each coefficient $(\mathbf{Z}\tilde{\mathbf{s}})_j$ for $j \in [m]$ is bounded above by $\alpha q n^2 d(\eta-1)$.

Proof: Using the same reasoning as in the proof of Lemma 1.13, we obtain $\sigma(\mathbf{Z}\tilde{\mathbf{s}}) = \theta(\mathbf{Z})\sigma(\tilde{\mathbf{s}})$ and thus $\|\sigma(\mathbf{Z}\tilde{\mathbf{s}})\|_2 \leq \|\theta(\mathbf{Z})\|_2 \cdot \|\sigma(\tilde{\mathbf{s}})\|_2$. Using Equation 1.1, i.e., the Vandermonde matrix \mathbf{V} to switch from the coefficient embedding τ to the canonical embedding σ , we obtain $\|\sigma(\tilde{\mathbf{s}})\|_2 \leq \|\mathbf{V}\|_2 \cdot \|\tau(\tilde{\mathbf{s}})\|_2 \leq n \cdot (\eta-1)\sqrt{nd}$, where we use that for cyclotomic number fields it yields $\|\mathbf{V}\|_2 \leq \|\mathbf{V}\|_F = \left(\sum_{j,k \in [n]} |\alpha_j^{k-1}|^2\right)^{1/2} \leq n$ (as all α_j are units) and that $\tau(\tilde{\mathbf{s}})$ is a vector of dimension nd with coefficients in $\{0, \dots, \eta-1\}$. Further, for each $j \in [m]$ and $k \in [d]$ it holds $\|\sigma(z_{jk})\|_2 \leq \alpha q \sqrt{n}$ with probability $1 - 2^{-\Omega(n)}$, using the Gaussian tail bound Lemma 1.6. Hence,

$$\|\theta(\mathbf{Z})\|_2 \leq \|\theta(\mathbf{Z})\|_F = \sqrt{\sum_{j \in [m]} \sum_{k \in [d]} \sum_{\ell \in [n]} |\sigma_\ell(z_{jk})|^2} = \sqrt{\sum_{j \in [m]} \sum_{k \in [d]} \|\sigma(z_{jk})\|_2^2} \leq \alpha q \sqrt{nd} \sqrt{m}.$$

Combining both bounds proves the claim. ■

We also need two other lemmata related to the inner product over $K_{\mathbb{R}}^d$ between a Gaussian vector and an arbitrary one. In particular, we use Lemma 1.16 in the proof of Lemma 2.5 (reducing extended M-LWE to binary secret M-LWE) in order to decompose a Gaussian noise into an inner product.

Lemma 1.15 ([LS15, Lem. 2.13])

Let $\mathbf{r} \in (\mathbb{R}^+)^n \cap H$, $\mathbf{z} \in K^d$ fixed and $\mathbf{e} \in K_{\mathbb{R}}^d$ sampled from $D_{\sqrt{\Sigma}}$, where $\sqrt{\Sigma} = [\delta_{j,k} \text{diag}(\mathbf{r})]_{j,k \in [d]} \in \mathbb{R}^{nd \times nd}$. Then $\langle \mathbf{z}, \mathbf{e} \rangle = \sum_{j \in [d]} z_j e_j$ is distributed according to $D_{\mathbf{r}'}$ with $r'_j = r_j \sqrt{\sum_{k \in [d]} |\sigma_j(z_k)|^2}$ for $j \in [n]$.

Lemma 1.16 (Adapted from [Reg09, Cor. 3.10])

Let $\mathcal{M} \subset K^d$ be an R -module (yielding a module lattice), let $\mathbf{u}, \mathbf{z} \in K^d$ be fixed, and let $\beta, \gamma > 0$. Assume that $(1/\beta^2 + \|\mathbf{z}\|_{2,\infty}^2/\gamma^2)^{-1/2} \geq \eta_\varepsilon(M)$ for some $\varepsilon \in (0, 1/2)$. Then the distribution of $\langle \mathbf{z}, \mathbf{v} \rangle + e$ where \mathbf{v} is sampled from $\mathcal{D}_{\mathcal{M}+\mathbf{u},\beta}$ and $e \in K_{\mathbb{R}}$ is sampled from D_γ , is within statistical distance at most 2ε from the elliptical Gaussian $D_{\mathbf{r}}$ over $K_{\mathbb{R}}$, where $r_j = \sqrt{\beta^2 \sum_{k \in [d]} |\sigma_j(z_k)|^2 + \gamma^2}$ for $j \in [n]$.

Proof: Consider $\mathbf{h} \in (K_{\mathbb{R}})^d$ distributed according to $D_{\mathbf{r}', \dots, \mathbf{r}'}$, where \mathbf{r}' is given by $r'_j = \gamma / \sqrt{\sum_{k \in [d]} |\sigma_j(z_k)|^2}$ for $j \in [n]$. Then by Lemma 1.15, $\langle \mathbf{z}, \mathbf{h} \rangle$ is distributed as D_γ and therefore $\Delta(\langle \mathbf{z}, \mathbf{v} \rangle + e, D_{\mathbf{r}}) = \Delta(\langle \mathbf{z}, \mathbf{v} + \mathbf{h} \rangle, D_{\mathbf{r}})$. Now, we denote \mathbf{t} such that $t_j = \sqrt{\beta^2 + (r'_j)^2}$ for $j \in [n]$. Note that by assumption

$$\begin{aligned} \min_{j \in [n]} \beta r'_j / t_j &= (1/\beta^2 + \max_{j \in [n]} \sum_{k \in [d]} |\sigma_j(z_k)|^2 / \gamma^2)^{-1/2} \\ &= (1/\beta^2 + \|\mathbf{z}\|_{2,\infty}^2 / \gamma^2)^{-1/2} \geq \eta_\varepsilon(M). \end{aligned}$$

Lemma 1.10 therefore applies and yields that $\mathbf{v} + \mathbf{h}$ is distributed as $D_{\mathbf{t}, \dots, \mathbf{t}}$, within statistical distance at most 2ε . By applying once more Lemma 1.15 and noticing that the statistical distance does not increase when applying a function (here the scalar product with \mathbf{z}), then we get that $\langle \mathbf{z}, \mathbf{v} + \mathbf{h} \rangle$ is distributed as $D_{\mathbf{r}}$ within statistical distance at most 2ε , where $r_j = t_j \sqrt{\sum_{k \in [d]} |\sigma_j(z_k)|^2} = \sqrt{\beta^2 \sum_{k \in [d]} |\sigma_j(z_k)|^2 + \gamma^2}$ for $j \in [n]$. ■

1.3.3 Statistical Distance

The statistical distance is an important measure of distance used in cryptography. We now recall its definition and some useful properties. Let P and Q be two discrete probability distributions over a discrete domain E . Their statistical distance is defined as

$$\Delta(P; Q) = \frac{1}{2} \sum_{x \in E} |P(x) - Q(x)|.$$

The statistical distance fulfills the probability preservation property.

Lemma 1.17 (Probability Preservation of the Statistical Distance)

Let P, Q be two probability distributions with $\text{Supp}(P) \subseteq \text{Supp}(Q)$ and $E \subseteq \text{Supp}(Q)$ be an arbitrary event. Then, $P(E) \leq \Delta(P; Q) + Q(E)$.

Within the thesis, we need the following two results about the statistical distance of two Gaussian distributions.

Lemma 1.18 ([DMR18, Thm. 1.2])

Let D_g denote the continuous Gaussian distribution over $K_{\mathbb{R}}$ and let $z \in K$. The statistical distance between D_g and $D_{g,z}$ is bounded above by $\frac{\sqrt{2\pi}\|z\|_2}{g}$.

Lemma 1.19 ([Reg05, Claim 2.2])

Let α and β be positive reals such that $\alpha < \beta$. The statistical distance between D_α and D_β is bounded above by $10 \cdot \left(\frac{\beta}{\alpha} - 1\right)$.

1.3.4 Rényi Divergence

The Rényi divergence [R61, vEH14] defines another measure of distribution closeness and was first used in cryptography as a powerful alternative for the statistical distance measure by Bai et al. [BLL⁺15, BLR⁺18]. In this thesis, it suffices to use the Rényi divergence of order 2. As in [BLR⁺18], we use the exponential of the standard definition. Let P and Q be two discrete probability distributions such that $\text{Supp}(P) \subset \text{Supp}(Q)$. The Rényi divergence of order 2 is defined as

$$\text{RD}_2(P\|Q) = \sum_{x \in \text{Supp}(P)} \frac{P(x)^2}{Q(x)} = E_{x \leftarrow P} \frac{\Pr[P = x]}{\Pr[Q = x]}.$$

From its definition, it directly follows for Q the uniform distribution over $\text{Supp}(P)$ that $\text{RD}_2(P\|Q) = |\text{Supp}(P)| \cdot \sum_{x \in \text{Supp}(P)} P(x)^2$, where $\sum_{x \in \text{Supp}(P)} P(x)^2$ is the collision probability of P . The Rényi divergence admits the following properties, proved in [vEH14].

Lemma 1.20 (Properties of the Rényi Divergence)

Let P, Q be two discrete probability distributions with $\text{Supp}(P) \subseteq \text{Supp}(Q)$. Further, let $(P_j)_j, (Q_j)_j$ be two families of independent discrete probability distributions with $\text{Supp}(P_j) \subset \text{Supp}(Q_j)$ for all j . Then, the following properties are fulfilled:

Data Processing Inequality: $\text{RD}_2(P^g\|Q^g) \leq \text{RD}_2(P\|Q)$ for any function g , where P^g (resp. Q^g) denotes the distribution of $g(y)$ induced by sampling $y \leftarrow P$ (resp. $y \leftarrow Q$),

Multiplicativity: $\text{RD}_2\left(\prod_j P_j \parallel \prod_j Q_j\right) = \prod_j \text{RD}_2(P_j\|Q_j)$,

Probability Preservation: Let $E \subset \text{Supp}(Q)$ be an arbitrary event, then

$$Q(E) \cdot \text{RD}_2(P\|Q) \geq P(E)^2.$$

In Chapter 2, we need the Rényi divergence of two shifted discrete Gaussians. By comparing the conditions of the lemma below for the Rényi divergence with the ones of Lemma 1.18 for the statistical distance, we can see that the first is a stronger measurement than the second. More precisely, a small Rényi divergence of two shifted Gaussians implies a small statistical distance.

Lemma 1.21 (Adapted from [LSS14, Lem. 4.2])

Let s and ε be positive real numbers with $\varepsilon \in (0, 1)$, \mathbf{c} be a vector of \mathbb{R}^n and Λ be a full-rank lattice in \mathbb{R}^n . We assume that $s \geq \eta_\varepsilon(\Lambda)$. Then,

$$\text{RD}_2(\mathcal{D}_{\Lambda, s, \mathbf{c}} \| \mathcal{D}_{\Lambda, s}) \leq \left(\frac{1 + \varepsilon}{1 - \varepsilon} \right)^2 \cdot \exp\left(\frac{2\pi \|\mathbf{c}\|_2^2}{s^2} \right).$$

1.3.5 Leftover Hash Lemma

The Leftover Hash Lemma (LHL) [HILL99] is a standard tool in cryptography used for security proofs. Within this thesis, we use two different flavors of it. The first one is a generalized version of the LHL in terms of universal hash functions, which we need when proving the security of our PKE scheme from Chapter 6. First, we recall the definition of a universal family of hash functions.

Definition 1.7 (Universal Family of Hash Functions): A (finite) family H of hash functions $h: X \rightarrow Y$ is called universal if

$$\Pr_{h \leftarrow U(H)} [h(x_1) = h(x_2)] = \frac{1}{|Y|},$$

for all $x_1 \neq x_2 \in X$.

We now recall the variant of the Leftover Hash Lemma as stated in [RSSS17].

Lemma 1.22 (Generalized Leftover Hash Lemma [RSSS17, Lem. 2.1])

Let X, Y and Z be finite sets, H be a universal family of hash functions $h: X \rightarrow Y$ and $g: X \rightarrow Z$ be an arbitrary function. Then, for any random variable T taking values in X we have

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

where $\gamma(T) = \max_{t \in X} \Pr [T = t]$.

The second variant of the LHL that we use within this work is an adaptation of the one by Micciancio [Mic07], which, instead of working with vectors over the finite field \mathbb{Z}_q , operates over principal ideal domains. Given a number field $K = \mathbb{Q}(\zeta)$, where the corresponding ring of

integers has the form $R = \mathbb{Z}[\zeta]$, and a prime q , then $R_q = R/qR$ is a principal ideal domain, allowing for a unique prime-ideal factorization. The result can easily be generalized to any number field $K = \mathbb{Q}(\zeta)$ with ring of integers R with $\mathbb{Z}[\zeta] \subset R$, as long as q is coprime with the index $[R : \mathbb{Z}[\zeta]]$. In this case, as demonstrated in [PP19, Lem. 2.14], there exists a bijection between $\mathbb{Z}[\zeta]$ and R .⁴ The only difference between the original statement and the one below is that the original proof uses the special number field $K \cong \mathbb{Q}[x]/\langle x^n - 1 \rangle$. Further, we state the LHL for a matrix-vector product instead of a vector-vector product and provide not only a bound on the statistical distance, but also on the Rényi divergence.

Lemma 1.23 (LHL over Rings adapted from [Mic07, Thm. 4.2])

Let $n, k, d, q, \eta \in \mathbb{N}$ such that q is prime. Further, let $K = \mathbb{Q}(\zeta)$ be a number field with associated ring of integers R such that q is coprime with the index $[R : \mathbb{Z}[\zeta]]$. We set $R_q = R/qR$ and $R_\eta = R/\eta R$. Then,

$$\Delta((\mathbf{C}, \mathbf{Cz}); (\mathbf{C}, \mathbf{s})) \leq \frac{1}{2} \sqrt{\left(1 + \frac{q^k}{\eta^d}\right)^n - 1}, \text{ and}$$

$$\text{RD}_2((\mathbf{C}, \mathbf{Cz}) \| (\mathbf{C}, \mathbf{s})) \leq \left(1 + \frac{q^k}{\eta^d}\right)^n,$$

where $\mathbf{C} \leftarrow U((R_q)^{k \times d})$, $\mathbf{z} \leftarrow U((R_\eta)^d)$ and $\mathbf{s} \leftarrow U((R_q)^k)$.

1.4 Learning With Errors

At the heart of most lattice-based schemes, and in particular at the core of this thesis, lies one computational problem, the Learning With Errors (LWE) problem. In this section, we introduce the problem, together with different structured variants of it, which are used in efficient schemes nowadays. We start with varying the underlying mathematical structure: First, we introduce the LWE problem over rational integers as it was originally presented by Regev [Reg05]. Then, we define LWE over rings of polynomials and finally over modules. Later in this section, we present two other versions of LWE, a first variant using the middle-product and a second one using deterministic rounding.

1.4.1 Learning With Errors (LWE)

The LWE problem, introduced by Regev [Reg05] in his pioneering work, serves as a fundamental computational problem in lattice-based cryptography. Informally speaking, an instance of the LWE problem is a system of noisy linear modular equations over the rational integers. Its search variant asks to find a solution to this system, whereas its decision version asks to distinguish such noisy linear modular equations from uniformly random ones. We now present the formal definitions of the LWE distribution and the LWE problem, in both its search and decision versions. Let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ denote the segment $[0, 1)$ with addition modulo 1.

Definition 1.8 (LWE Distribution): Let $q, n \in \mathbb{N}$ such that $q \geq 2$. Further, let ψ denote a distribution over \mathbb{T} and fix a vector $\mathbf{s} \in \mathbb{Z}_q^n$. We set $A_{\mathbf{s}, \psi}$ as the distribution over $\mathbb{Z}_q^n \times \mathbb{T}$

⁴As we only became aware of this bijection recently, the original results [BJRW20, BJR21] restrict the LHL to number fields K , where R_q is a principal ideal domain. Furthermore, we state a more general version allowing the secret to have larger coefficients.

obtained by choosing a vector $\mathbf{a} \leftarrow U(\mathbb{Z}_q^n)$, an element $e \leftarrow \psi$ and returning $(\mathbf{a}, \frac{1}{q}\langle \mathbf{a}, \mathbf{s} \rangle + e \bmod \mathbb{Z})$.

Definition 1.9 (LWE): The search version $\text{SLWE}_{n,q,\psi}$ of the Learning With Errors problem is as follows: let $\mathbf{s} \in \mathbb{Z}_q^n$ be secret; given arbitrarily many samples from $A_{\mathbf{s},\psi}$, the goal is to find \mathbf{s} . Its decision version $\text{LWE}_{n,q,\psi}$ is as follows: choose $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$; the goal is to distinguish between arbitrarily many independent samples from $A_{\mathbf{s},\psi}$ and the same number of independent samples from $U(\mathbb{Z}_q^n \times \mathbb{T})$.

Sometimes, it is convenient to interpret LWE in terms of matrices. Let $m \in \mathbb{N}$ be the number of requested samples of the LWE distribution $A_{\mathbf{s},\psi}$, i.e., we are given $(\mathbf{a}_j, \frac{1}{q}\langle \mathbf{a}_j, \mathbf{s} \rangle + e_j \bmod \mathbb{Z})$ for $j \in [m]$. Then, we consider the matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ whose rows are the \mathbf{a}_j 's and we set $\mathbf{e} = (e_1, \dots, e_m)^T$. We obtain the matrix representation $(\mathbf{A}, \frac{1}{q}\mathbf{A} \cdot \mathbf{s} + \mathbf{e})$, where $\mathbf{s} \in \mathbb{Z}_q^n$. Note that LWE with respect to the matrix \mathbf{A} defines an instance of the computational problem BDD over the lattice $\Lambda_q(\mathbf{A}^T)$, as defined in Section 1.2.1. The promise bound needed for BDD is characterized by the error distribution ψ .

As an attractive property for cryptography, LWE enjoys quantum [Reg05, Reg09, PRS17a] and classical [Pei09, BLP⁺13] worst-case to average-case reductions for suitable parameter choices, from well-studied problems such as finding a set of short independent vectors (SIVP) or the decision variant of finding short vectors (GapSVP, Def 1.2) in Euclidean lattices. The term average-case is used to stress that LWE is defined over the randomness of \mathbf{A}, \mathbf{s} and \mathbf{e} . In other words, an algorithm that efficiently solves LWE can be used to classically (or quantumly) solve the problems SIVP or GapSVP in *any* Euclidean lattice. By classical algorithm we denote algorithms that use classical mechanics, whereas quantum algorithms operate with quantum mechanics. In order to perform a quantum algorithm, a so-called quantum computer is needed. A standard conjecture is to assume that there is no polynomial-time algorithm that solves these lattice problems and their approximated versions, even on quantum computers (see Conjecture 1 from the introduction). Thus, any solver of the average-case problem LWE can be transformed into a solver for any instance of the corresponding worst-case problem, which is presumed to be difficult.

For efficiency reasons, we are interested in variants of LWE, where the underlying secret has *small* norm. A relatively simple, but quite elegant reduction allows the secret to follow the same distribution as the noise [ACPS09]. This variant is commonly called the Hermite Normal Form (HNF) of LWE. By re-randomizing the secret we also obtain the reduction in the opposite direction, making both problems equally hard. Due to its widespread use in cryptographic constructions, and the equivalence of LWE and LWE in HNF, they are sometimes presented as the same problem.

1.4.2 Polynomial Learning With Errors (P-LWE)

The cryptographic protocols relying on the hardness of LWE are inherently inefficient due to the size of the public keys which usually contain the matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, where n is at least as large as the security parameter and m is the number of samples which is usually larger than $n \log(n)$. To improve the efficiency, structured variants of LWE have been proposed [SSTX09, LPR10, LS15]. One prominent variant is the Polynomial Learning With Errors (P-LWE) problem, introduced by Stehlé et al. [SSTX09]. Instead of considering noisy linear equations over the rational integers \mathbb{Z} , the linear equations are now defined algebraically over some ring over polynomials with

coefficients in \mathbb{Z} . As we only need the decision version throughout this work, we focus on it in the following, omitting the search variant.

Definition 1.10 (Decision P-LWE): Let $q, m \in \mathbb{N}$ such that $q \geq 2$. Let $f(x)$ be a monic and irreducible polynomial in $\mathbb{Z}[x]$ of degree m and ψ be a distribution over $\mathbb{R}[x]/\langle f(x) \rangle$. The decision P-LWE $_{q,\psi}^f$ problem asks to distinguish arbitrarily many samples of the form $(a_j, b_j = a_j \cdot s + e_j \bmod q)$, where $e_j \leftarrow \psi$ and $a_j \leftarrow U(\mathbb{Z}_q[x]/\langle f(x) \rangle)$, from the same number of samples chosen uniformly over $\mathbb{Z}_q[x]/\langle f(x) \rangle \times \mathbb{R}_q[x]/\langle f(x) \rangle$ with non-negligible success probability over the choices of $s \leftarrow U(\mathbb{Z}_q[x]/\langle f(x) \rangle)$.

As for LWE, the P-LWE problem also admits worst-case to average-case quantum reductions from well-studied lattice problems [SSTX09]. Whereas the hardness reductions for LWE start from lattice problems in the class of general Euclidean lattices, the class has to be restricted to *ideal lattices* in the case of P-LWE. As explained in Section 1.2.2, these ideal lattices correspond to the ideals in the polynomial ring $\mathbb{Z}[x]/\langle f(x) \rangle$.

1.4.3 Module Learning With Errors (M-LWE)

The module variant of LWE was first defined by Brakerski et al. [BGV12] and thoroughly studied by Langlois and Stehlé [LS15]. Instead of working over a ring of polynomials as for P-LWE, the module variant is defined over the ring of integers of some number field K . Since its introduction, the M-LWE problem has enjoyed more and more popularity as it offers a fine-grained trade-off between concrete security and efficiency. Within the NIST standardization process, several third round candidates rely on the hardness of M-LWE, e.g., the signature scheme **Dilithium** [DKL⁺18] and the key encapsulation mechanism **Kyber** [BDK⁺18]. We now present the formal definitions of the M-LWE distribution and the M-LWE problem, in both its search and decision versions. By \mathbb{T}_{R^\vee} we denote the torus $K_{\mathbb{R}}/R^\vee$, where $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$.

Definition 1.11 (M-LWE Distribution): Let K be a number field of degree n and R its ring of integers with dual R^\vee . Further, let $d \in \mathbb{N}$ denote the rank, let ψ be a distribution over $K_{\mathbb{R}}$ and let $\mathbf{s} \in (R_q^\vee)^d$ be a vector. We let $A_{\mathbf{s},\psi}^{\mathcal{M}}$ denote the distribution over $(R_q)^d \times \mathbb{T}_{R^\vee}$ obtained by choosing a vector $\mathbf{a} \leftarrow U((R_q)^d)$, an element $e \leftarrow \psi$ and returning $(\mathbf{a}, \frac{1}{q}\langle \mathbf{a}, \mathbf{s} \rangle + e \bmod R^\vee)$.

The inner product of the vectors $\mathbf{s} = (s_1, \dots, s_d)^T$ and $\mathbf{a} = (a_1, \dots, a_d)^T$ is defined in the natural way, i.e., $\langle \mathbf{a}, \mathbf{s} \rangle = \sum_{j \in [d]} a_j \cdot s_j$, using the addition and multiplication of the field K .

Definition 1.12 (M-LWE): Let $q, d \in \mathbb{N}$ such that $q \geq 2$. Let Ψ be a family of distributions over $K_{\mathbb{R}}$. The search version M-SLWE $_{n,d,q,\Psi}$ of the Module Learning With Errors problem is as follows: let $\mathbf{s} \in (R_q^\vee)^d$ be secret and $\psi \in \Psi$; given arbitrarily many samples from $A_{\mathbf{s},\psi}^{\mathcal{M}}$, the goal is to find \mathbf{s} . Let Υ be a distribution on a family of distributions over $K_{\mathbb{R}}$. Its decision version M-LWE $_{n,d,q,\Upsilon}$ is as follows: choose $\mathbf{s} \leftarrow U((R_q^\vee)^d)$ and $\psi \leftarrow \Upsilon$; the goal is to distinguish between arbitrarily many independent samples from $A_{\mathbf{s},\psi}^{\mathcal{M}}$ and the same number of independent samples from $U((R_q)^d \times \mathbb{T}_{R^\vee})$.

Module LWE as a Generalization

Historically, the M-LWE problem was introduced after another important variant of LWE, namely the Ring Learning With Errors (R-LWE) problem, as defined by Lyubashevsky et al. [LPR10]. The latter one can be seen as a special case of the first by specifying $d = 1$. The R-LWE problem coincides with the P-LWE problem from Section 1.4.2 for the class of cyclotomic fields (where ζ is a primitive root of the unity). Ro ca et al. [RSW18] show that P-LWE and R-LWE are equivalent (with some parameter losses) for a larger class of polynomials. In a more recent work, the hardness connection between M-LWE and R-LWE is further investigated. More precisely, Peikert and Pepin [PP19] show a tight reduction from R-LWE over a number field of degree $n \cdot k$ to M-LWE over a number field of degree n and with rank k . Furthermore, the special case of M-LWE, where the ring has degree $n = 1$ and thus $R = \mathbb{Z}$, is exactly the original LWE problem from Section 1.4.1. Thus, M-LWE can be seen as a generalization, encompassing both the unstructured LWE and the fully structured R-LWE (and hence P-LWE) problems.

Primal versus Dual

Note that we use the definition of M-LWE in its so-called *dual* form as the secret vector is taken over the dual R^\vee . There exists also a flavor of M-LWE where the secret vector is chosen over the ring R directly, referred to as M-LWE in its *primal* version. Ro ca et al. [RSW18] show for R-LWE that both variants are equivalent up to some polynomial losses in the error and their results can easily be adapted to the module setting. The reason to work with the dual is that it naturally arises in the worst-to-average case reduction and thus leads to tighter security proofs [LPR10, Sec. 3.3]. Furthermore, it allows for optimal lattice decoding algorithms as explained in [LPR13].

Worst-Case Reductions

Similar to its unstructured counterpart, M-LWE also enjoys worst-case to average-case connections for suitable parameter choices from lattice problems such as SIVP $_\gamma$ [LS15]. Whereas the hardness results for LWE start from lattice problems in the class of general Euclidean lattices, the set has to be restricted to *module lattices* in the case of M-LWE. As introduced in Section 1.2.2, these module lattices correspond to modules in the ring R and we refer to the related lattice problem as Mod-SIVP $_\gamma$ and Mod-GapSVP $_\gamma$ (Def. 1.2), respectively. As an additional feature, a converse reduction from M-LWE to Mod-SIVP $_\gamma$ is proven for the special case of power-of-2 cyclotomics [LS15] and improved by Wang and Wang [WW19] for all cyclotomic fields.

Matrix Representation

As for LWE, we can interpret the M-LWE distribution in terms of matrices. Let $m \in \mathbb{N}$ be the number of requested samples of the M-LWE distribution $A_{\mathbf{s}, \psi}^m$, given by $(\mathbf{a}_j, \frac{1}{q} \langle \mathbf{a}_j, \mathbf{s} \rangle + e_j)$ for $j \in [m]$. Then, we consider the matrix $\mathbf{A} \in R_q^{m \times d}$ whose rows are the \mathbf{a}_j 's and we set $\mathbf{e} = (e_1, \dots, e_m)^T$. We obtain the representation $(\mathbf{A}, \frac{1}{q} \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$, where $\mathbf{s} \in (R_q^\vee)^d$. By adding the parameter m as a superscript suffix, we denote this problem by M-LWE $_{n,d,q,\Upsilon}^m$. Furthermore, we can interpret the matrix $\mathbf{A} \in R_q^{m \times d}$ as an $nm \times nd$ matrix over \mathbb{Z} , where n is the ring degree of R . As explained in Section 1.1.3, the multiplication of a ring element $a \in R_q$ with another ring element $s \in R_q$ (with respect to their coefficient embeddings) can be expressed as the matrix-vector product using the matrix $\text{Rot}(a) \in \mathbb{Z}_q^{n \times n}$. Thus, the $m \times d$ matrix $\mathbf{A} = (a_{jk})_{jk}$ over R_q defines an $nm \times nd$ matrix $\text{Rot}(\mathbf{A}) = (\text{Rot}(a_{jk}))_{jk}$ over \mathbb{Z}_q . In other words, an instance of the M-LWE distribution can be viewed as an instance of the LWE distribution, where the public matrix $\text{Rot}(\mathbf{A})$ is not fully

random, but has some structure. This is why we often refer to it as a *structured* variant of LWE, see Figure 4 from the introduction for an illustration.

Multiple Secrets

Another variant of M-LWE is the multiple secrets M-LWE problem. Let $k, m \in \mathbb{N}$, where m denotes the number of requested samples of $A_{\mathbf{s}, \psi}^M$, and k stands for the number of independent secrets. In the multiple secrets version, the secret vector $\mathbf{s} \in (R_q^\vee)^d$ is replaced by a secret matrix $\mathbf{S} \in (R_q^\vee)^{d \times k}$ and the error vector $\mathbf{e} \leftarrow \psi^m$ by an error matrix $\mathbf{E} \leftarrow \psi^{m \times k}$. There is a simple polynomial-time reduction from M-LWE using a secret vector to M-LWE using a secret matrix for any k polynomially large in the module rank d via a hybrid argument, as given for instance in [Mic18, Lem. 2.9]. By adding the parameter k as a superscript suffix, we denote the corresponding problem by $\text{M-LWE}_{n,d,q,\Upsilon}^{m,k}$.

Small Secrets

Another possibility is to choose a *small* secret. As for LWE, we can define the Hermite Normal Form (HNF) of M-LWE, where the secret follows the same distribution as the noise, and show that both problems are equally hard by adapting the proof for LWE [ACPS09]. The HNF of M-LWE is often used in cryptographic constructions to improve efficiency, as for instance in signature schemes like Dilithium [DKL⁺18]. Alternatively, we can define a variant of M-LWE, where the coefficients of the secret vector are within a ball of small radius. More precisely, for some positive integer η we select uniformly at random the coefficients of the secret vector over the set R_η^\vee , where $R_\eta = \{x \in R: \tau(x) \in \{0, \dots, \eta - 1\}^n\}$. We denote the corresponding problem by η -M-LWE. In Chapter 2, we are interested in the special case of η -M-LWE, where the secret vector \mathbf{s} is a binary vector, thus chosen over the set R_2^\vee (i.e., $\eta = 2$). As this variant plays an important role within this thesis, we separately denote it by bin-M-LWE . As we argue later in Section 2.5, the set R_η^\vee is defined with respect to the coefficient embedding τ .

Discrete Variant

As pointed out by Lyubashevsky et al. [LPR10], sometimes it can be convenient to work with a discrete variant, where the second component b of each sample (\mathbf{a}, b) is taken from a finite set, and not from the continuous domain \mathbb{T}_{R^\vee} . Indeed, for the case of M-LWE, if the rounding function $\lfloor \cdot \rfloor : K_{\mathbb{R}} \rightarrow R^\vee$ is chosen in a suitable way, see Lyubashevsky et al. [LPR13, Sec. 2.6] for more details, then every sample $(\mathbf{a}, b = \frac{1}{q} \langle \mathbf{a}, \mathbf{s} \rangle + e) \in (R_q)^d \times \mathbb{T}_{R^\vee}$ of the distribution $A_{\mathbf{s}, \psi}^M$ can be transformed to $(\mathbf{a}, \lfloor q \cdot b \rfloor) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + \lfloor q \cdot e \rfloor \bmod qR^\vee) \in (R_q)^d \times R_q^\vee$. For technical reasons, we use the latter representation in Section 2.2.

1.4.4 Middle-Product Learning With Errors (MP-LWE)

We now present another variant of LWE which can be seen as a modification of P-LWE, where the standard product of two polynomials is replaced by their so-called middle-product. The use of the middle-product in lattice-based cryptography was introduced by Roca et al. [RSSS17] in the form of Middle-Product Learning With Errors (MP-LWE). We first recall the definition of the middle-product of two polynomials of bounded degree. We denote by $\mathbb{Z}_q^{<n}[x]$ the set of polynomials in $\mathbb{Z}_q[x]$ with degree less than n .

where $\text{rev}(\mathbf{b})$ and $\text{rev}(s)$ denote the coefficient vectors of b and s in reverse order.

We now formally define the Middle-Product Learning With Errors (MP-LWE) problem. As we only need the decision version throughout this work, we focus on it in the following.

Definition 1.15 (Decision MP-LWE): Let $q, d, n \in \mathbb{N}$ such that $q \geq 2$ and $0 < d \leq n$. Further, let ψ be a distribution over $\mathbb{R}^{<d}[x]$. The decision MP-LWE $_{q,n,d,\psi}$ problem asks to distinguish arbitrary many samples of the form $(a_j, b_j = a_j \odot_d s + e_j \bmod q)$, where $e_j \leftarrow \psi$ and $a_j \leftarrow U(\mathbb{Z}_q^{<n}[x])$, from the same number of samples chosen uniformly over $\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x]$ with non-negligible success probability over the choices of $s \leftarrow U(\mathbb{Z}_q^{<n+d-1}[x])$.

Roşca et al. [RSSS17] show that MP-LWE is at least as hard as P-LWE (Def. 1.10) for many polynomials $f(x) \in \mathbb{Z}_q[x]$. The noise parameter of the MP-LWE problem can be set to handle an exponentially large class of polynomials $f(x)$. A more recent result by Peikert and Pepin [PP19] shows a tighter (and more direct) reduction from R-LWE to MP-LWE. In both settings, the definition of MP-LWE is independent of some particular defining polynomial $f(x)$ and at the same time its hardness is for suitable parameters still implied by worst-case problems over ideal lattices.

1.4.5 Learning With Rounding (LWR)

In 2012, Banerjee et al. [BPR12] introduce a deterministic variant of LWE, namely the Learning With Rounding (LWR) problem. Originally, it is used to construct efficient pseudorandom functions [BPR12], lossy trapdoor functions and deterministic encryption schemes [AKPW13]. Structured variants of LWR also serve as the underlying hardness assumption for efficient public key encryption schemes [DKRV18, BBF⁺19], currently considered for standardization, and efficient pseudorandom functions [CS19].

Before we formally present this deterministic variant of LWE, we need to define a method to round elements from \mathbb{Z}_q to \mathbb{Z}_p , where $p \leq q$.

Definition 1.16 (Rounding): Let $p, q \in \mathbb{N}$ such that $2 \leq p \leq q$. The modular rounding function $\lfloor \cdot \rfloor_p : \mathbb{Z}_q \rightarrow \mathbb{Z}_p$ is defined as $\lfloor x \rfloor_p = \left\lfloor \left(\frac{p}{q} \right) \cdot x \right\rfloor \bmod p$, where $\lfloor \cdot \rfloor$ is the standard rounding function, mapping every $y \in \mathbb{R}$ to its closest integer $\lfloor y \rfloor \in \mathbb{Z}$.

The modular rounding function extends component-wise to vectors over \mathbb{Z}_q and coefficient-wise to polynomials in $\mathbb{Z}_q[x]$. Note that we use the same notation as Banerjee et al. [BPR12] for the purpose of coherence. It is also possible to use the floor rounding function $\lfloor \cdot \rfloor$, where each element is rounded down to the next smaller integer, as for instance done by Chen et al. [CZZ18].

In order to lift rounded elements from \mathbb{Z}_p back to \mathbb{Z}_q , we define a probabilistic lifting function $\text{lift}(\cdot) : \mathbb{Z}_p \rightarrow \mathbb{Z}_q$ that takes $x \in \mathbb{Z}_p$ as input and chooses uniformly at random an element u from the set $\{v \in \mathbb{Z}_q : \lfloor v \rfloor_p = x\}$. Again, $\text{lift}(\cdot)$ can be extended coefficient-wise to vectors over \mathbb{Z}_q and polynomials in $\mathbb{Z}_q[x]$. This lifting function becomes important in the encryption scheme in Chapter 6. There, we use $\text{lift}(\cdot)$ to lift rounded polynomials in $\mathbb{Z}_p[x]$ back to $\mathbb{Z}_q[x]$ such that $\left\lfloor \text{lift}(\lfloor a \rfloor_p) \right\rfloor_p = \lfloor a \rfloor_p$. Note that $\text{lift}(\lfloor a \rfloor_p) = a + e$ with $\|e\|_\infty \leq \frac{q}{p}$. Alternatively, one can

deterministically lift any element from \mathbb{Z}_p to \mathbb{Z}_q by mapping $y \in \mathbb{Z}_p$ to $\lfloor q/p \cdot y \rfloor \in \mathbb{Z}_q$. This induces a rounding error whose norm is bounded above by q/p , as before with the probabilistic lifting function $\text{lift}(\cdot)$.

We now define the Learning With Rounding problem. Again, we focus on its decision version throughout the thesis.

Definition 1.17 (Decision LWR): Let $q, p, n \in \mathbb{N}$ such that $2 \leq p \leq q$. The decision Learning With Rounding (LWR) problem ask to distinguish arbitrary many samples of the form $(\mathbf{a}_j, b_j = \lfloor \langle \mathbf{a}_j, \mathbf{s} \rangle \rfloor_p)$, where $\mathbf{a}_j \leftarrow U(\mathbb{Z}_q^n)$, from the same number of samples chosen uniformly over $\mathbb{Z}_q^n \times \mathbb{Z}_q$ with non-negligible success probability over the choices of $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$.

Hardness of LWR

When first introducing the problem LWR, Banerjee et al. [BPR12] show a reduction from LWE to LWR with arbitrarily many samples. 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 propose new reductions that work for a larger range of parameters. Alwen et al. [AKPW13] give 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 a priori. Further, the reduction imposes certain number theoretical restrictions on the modulus q . For example, power-of-2 moduli are not covered. In a more recent work, Bogdanov et al. [BGM⁺16] use the Rényi divergence (see Section 1.3.4) to show a sample preserving reduction for the search variant. Moreover, they provide a reduction for the decision variant using Fourier learning techniques, which is dimension preserving for the special case that the modulus is prime. It is slightly improved to composite moduli by Bai et al. [BLL⁺15]. In a concurrent work, Alperin-Sheriff and Apon [AA16] further improve the parameter sets for the reduction. In particular, the reduction is dimension-preserving for any polynomial-sized modulus.

Structured Variants of LWR

As for LWE, the definition of LWR can be adapted to the ring setting, defining the module variant M-LWR and the ring variant R-LWR. When defining M-LWR and R-LWR, one has to decide how to round vectors over R_q . One possibility is to consider an element of R_q via its coefficient embedding as a vector over \mathbb{Z}_q and simply round each coefficient to \mathbb{Z}_p . This is the way chosen by most works considering structured variants of LWR, as for instance [BPR12, BGM⁺16, AA16, CZZ18]. In contrast, a more recent work by Liu and Wang [LW20] studies a variant of R-LWR where the rounding is done with respect to so-called normal integral bases.

Concerning the hardness of R-LWR, in the full version of their paper, published on the IACR eprint server, Banerjee et al. [BPR11] show that their reduction with arbitrarily many samples also works for the ring counterpart. However, as for LWR, this reduction similarly imposes a super-polynomial modulus. The improved reduction of Bogdanov et al. [BGM⁺16] also applies to the search variant of R-LWR, but not to the decision variant (as they use the Rényi divergence). In the same manner, some parts of the reduction of Alperin-Sheriff and Apon [AA16] also apply to R-LWR. All those works above leave it as an open problem to prove the hardness of R-LWR

(or M-LWE) in its decision variant, which is usually needed to construct provably secure encryption schemes. To overcome this issue, Chen et al. [CZZ18] propose a new (non-standard) assumption, called the Computational Learning With Rounding Over Rings (R-CLWR) problem. They show an efficient reduction from decision R-LWE to R-CLWR, and construct an practical encryption scheme based on the hardness of R-CLWR. Recently, Liu and Wan [LW20] address the remaining open problem by providing a search-to-decision reduction for R-LWR. In order to do so, they need to define a new way of rounding with respect to normal integral bases.

1.4.6 Practical Hardness of Learning With Errors

In this thesis, we are mainly studying the *theoretical* hardness of LWE (and its variants) by providing efficient reductions from a well studied problem and thus preserving the worst-case to average-case connections that make LWE so interesting to use in cryptography. However, if one constructs a cryptographic scheme whose security is based on LWE (or some of its variants), deriving concrete parameters from those reductions would lead to rather inefficient schemes. This is why in practice, it is common to derive the concrete parameters for LWE-based schemes by looking at the cost of the best known algorithms solving LWE. We call this the *practical* hardness of LWE.

Fortunately, there are handy tools that allow to estimate the practical hardness of LWE, depending on its different parameters. Using for instance the LWE Estimator by Albrecht et al. [APS15] publicly available at <https://bitbucket.org/malb/lwe-estimator/src>, one can measure the practical hardness of LWE. The more recent Leaky Estimator by Dachman-Soled et al. [DDGR20] publicly available at <https://github.com/lucas/leaky-LWE-Estimator> also allows to integrate different types of hints on the LWE secret and/or noise. The latter is of important use to refine the security analysis of the encryption scheme presented in Chapter 7.

1.5 Cryptographic Notions

In this section, we introduce the notion of Public Key Encryption schemes. For a gentle introduction to modern cryptography we refer to the standard book of Katz and Lindell [KL14].

1.5.1 Public Key Encryption

A Public-Key Encryption (PKE) scheme permits two parties to confidentially exchange messages without sharing a common secret key beforehand. In the following, we provide formal definitions of PKE schemes, their correctness and IND-CPA security properties.

Definition 1.18 (Public Key Encryption): A Public Key Encryption (PKE) scheme $\Pi = (\text{KGen}, \text{Enc}, \text{Dec})$ for a message space M and a ciphertext space C is composed of three PPT algorithms, specified as follows:

KGen: The key generation algorithm KGen takes as input a security parameter λ and returns a key pair (sk, pk) , called the secret key sk and the public key pk .

Enc: The encryption algorithm Enc takes as input the public key pk and a message $m \in M$ and returns the ciphertext $c \leftarrow \text{Enc}(\text{pk}, m) \in C$.

Dec: The decryption algorithm Dec takes as input the secret key sk and a ciphertext $c \in C$ and returns $\text{Dec}(\text{sk}, c) \in M \cup \{\perp\}$, where \perp denotes the failure symbol.

Definition 1.19 (Correctness): We call the PKE scheme $\Pi = (\text{KGen}, \text{Enc}, \text{Dec})$ correct with correctness error $\delta \in [0, 1)$ if for any message $m \in \mathcal{M}$ it yields

$$\Pr[\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m)) = m] \geq 1 - \delta,$$

where the probability is taken over the key pair $(\text{sk}, \text{pk}) \leftarrow \text{KGen}(1^\lambda)$ and the randomness used by the encryption and decryption algorithms. If $\delta = 0$, we say that Π is perfectly correct.

Informally speaking, the security notion IND-CPA captures that no efficient adversary can distinguish between the ciphertext of two messages, where the adversary has even the right to choose the messages by themselves. The acronym stands for *indistinguishable against chosen-plaintext attacks*.

Definition 1.20 (IND-CPA Security): We say that the PKE scheme $\Pi = (\text{KGen}, \text{Enc}, \text{Dec})$ is IND-CPA secure, if for all PPT adversaries \mathcal{A} , there is a negligible function $\text{negl}(\cdot)$ such that

$$\Pr[\text{IND-CPA}_{\Pi}^{\mathcal{A}} = 1] < \frac{1}{2} + \text{negl}(\lambda),$$

where $\text{IND-CPA}_{\Pi}^{\mathcal{A}}$ is the security game from Protocol 1.1.

Protocol 1.1: The IND-CPA security game.

IND-CPA _Π ^ℳ
1: $b \leftarrow U(\{0, 1\})$
2: $(\text{pk}, \text{sk}) \leftarrow \text{KGen}(1^\lambda)$
3: $(m_0, m_1) \leftarrow \mathcal{A}(1^\lambda, \text{pk})$
4: $c \leftarrow \text{Enc}(\text{pk}, m_b)$
5: $b' \leftarrow \mathcal{A}(1^\lambda, \text{pk}, c)$
6: return $b = b'$

1.5.2 Random Oracle Model

Within this thesis, we make use of the Random Oracle Model (ROM). This is an idealized framework, where we assume the existence of perfectly random functions, realized by oracles. More precisely, for $n, k \in \mathbb{N}$, a random oracle $H: \{0, 1\}^k \rightarrow \{0, 1\}^n$, queried on some input bit string x of length k outputs a truly random bit string $H(x)$ of length n . Everyone can interact with the random oracle, but how it internally works stays unknown to everybody. We require the random oracle to be consistent, i.e., if it already has been queried on some input x , its answer is consistent with the previous one. It can be useful to imagine that the random oracle maintains a list, which is initially empty and where input-output pairs are added on-the-fly. A useful property for security proofs is that we can only guess the answer of H on a given input x as long as we don't query it.

We say that a security proof is in the ROM if it makes use of random oracles. The main advantage of such proofs is that they offer provably secure cryptographic schemes that are in

most cases more efficient than their analogues proven in the standard model. As we don't know any realization of random oracles in the real work, when implementing such schemes, the random oracles are instantiated by concrete cryptographic hash functions. Even as there is no theoretical evidence that a security proof in the ROM implies the security of a scheme when the random oracle is replaced by a specific cryptographic hash function, it still increases our confidence in the scheme's security and can be seen as a compromise between efficient schemes without security proof and impractical schemes with security proof. For a more detailed discussion we refer to the introduction of Katz and Lindell [KL14].

Part I

Theoretical Foundations

Chapter 2

Binary Hardness of Module LWE

The content of this chapter is based on two joint works with Corentin Jeudy, Adeline Roux-Langlois and Weiqiang Wen. The first is published in the proceedings of the conference Asiacrypt 2020 [BJRW20] and the second is published in the proceedings of the conference CT-RSA 2021 [BJRW21]. For the latter, a presentation of 20 minutes given by one of the co-authors has been recorded for the conference, illustrating the high level techniques of the paper.¹

Contents

2.1	Introduction	33
2.1.1	Our Contributions	34
2.1.2	Related Work	36
2.1.3	Roadmap	38
2.2	Warm-up: A Simple Reduction	38
2.3	An Improved Reduction	41
2.3.1	First-is-errorless M-LWE	44
2.3.2	Extended M-LWE	45
2.3.3	Reduction to Binary M-LWE	50
2.4	Generalization to Larger Secrets	53
2.5	Choice of Embedding for Binary Secrets	55
2.5.1	Lagrange Basis	55
2.5.2	Lagrange Basis for Power-of-2 Cyclotomics	56

2.1 Introduction

The Learning With Errors (LWE) problem, as introduced in Section 1.4, is one of the fundamental problems used in lattice-based cryptography. Since its introduction in 2005 by Regev [Reg05], several variants of it have been studied. One line of research investigates the hardness of LWE when the underlying secret is of *small* norm. A relatively simple, but elegant reduction allows the secret to follow the same distribution as the noise [ACPS09]. This variant is commonly called the Hermite Normal Form of LWE. Another, more extreme case is to choose the secret as a binary vector, resulting in an LWE sample, where $\mathbf{s} \in \{0, 1\}^n$ instead of $\mathbf{s} \in \mathbb{Z}_q^n$. We call this the binary secret LWE problem, denoted by bin-LWE. It is particularly interesting as it increases efficiency.

¹<https://www.youtube.com/watch?v=AfDskDKEzvg> starting from 23 min 35 sec.

Furthermore, modulus-rank switching techniques [BLP⁺13, AD17a, WW19] rely on using small secrets as they keep the noise blowup to a minimum. The binary secret variant also happens to be essential for some Fully Homomorphic Encryption (FHE) schemes as in [DM15, CGGI16], as it helps to control the noise growth during computations on encrypted data.

A first study of bin-LWE is provided by Goldwasser et al. [GKPV10] in the context of leakage-resilient cryptography. Whereas their proof structure has the advantage of being easy to follow, their result suffers from a huge error increase. Informally, they show a reduction from $\text{LWE}_{k,q,D_\alpha}$ to $\text{bin-LWE}_{n,q,D_\beta}$, where $\frac{\alpha}{\beta} = \text{negl}(n)$ and $n \geq k \log_2 q + \omega(\log_2 n)$. Recall, that D_α denotes the continuous Gaussian distribution of width α . Later, Brakerski et al. [BLP⁺13] improve the state of the art in order to show the classical hardness of LWE with a polynomial-sized modulus. Micciancio [Mic18] provides another reduction from LWE to its binary version. Whereas the two reduction techniques differ, both papers achieve similar results. The dimension is still increased roughly by a factor $\log_2 q$, but the error only by a factor of \sqrt{n} , where n is the resulting LWE dimension. More concretely, in [BLP⁺13] a reduction from $\text{LWE}_{k,q,D_\alpha}$ to $\text{bin-LWE}_{n,q,D_\beta}$ is shown, where $\frac{\alpha}{\beta} \leq \frac{1}{\sqrt{10n}}$ and $n \geq (k+1) \log_2 q + \omega(\log_2 n)$. And Micciancio [Mic18] proves a reduction from $\text{LWE}_{k,q,D_\alpha}$ to $\text{bin-LWE}_{n,q,D_\beta}$, where $\frac{\alpha}{\beta} \leq \frac{1}{2\sqrt{n+1}}$ and $n \geq (k+1) \log_2 q + \omega(\log_2 n)$. The increase in dimension from k to roughly $k \log_2 q$ is reasonable, as it essentially preserves the number of possible secrets. However, all of the proofs mentioned above work over the integers \mathbb{Z} and not over the ring of integers R of some number field. Hence, as stated by Micciancio in the conclusion of [Mic18], an important open problem is whether similar results carry over to the structured variants, in particular to M-LWE, as practical schemes are based rather on M-LWE than on LWE for efficiency reasons.

2.1.1 Our Contributions

In this chapter, we prove the hardness of the binary secret version of M-LWE, if the module rank d is (super-)logarithmic in the degree n of the underlying number field. To the best of our knowledge, this is the first result on the hardness of structured LWE with binary secrets.

More precisely, we present two different proofs to obtain this result. The first proof, presented in Section 2.2, can be seen as a warm-up for the second one, subsequently presented in Section 2.3. Whereas the latter obtains better noise parameters than the first, it is also more complex and requires the introduction of several intermediate problems, such as extended M-LWE and first-is-errorless M-LWE. From a high level perspective, the first more simple proof follows the techniques of [GKPV10], while improving the noise rate significantly by using the Rényi divergence (RD) instead of the statistical distance. The second more involved proof can be seen as a module version of [BLP⁺13], where moving to the module setting of LWE imposes some technical difficulties that we need to address carefully. We summarize the different characteristics of both reductions in Table 2.1. The most remarkable difference between them is that the simple reduction, represented in the left column, allows for a rank d that is logarithmic in the ring degree n (i.e., $\Omega(\log_2 n)$), whereas the improved reduction, stated in the right column, only allows for a rank d that is super-logarithmic in n (i.e., $\omega(\log_2 n)$). In contrast, the latter improves the noise ratio of the first one by roughly a factor of \sqrt{md} , where m denotes the number of given samples.

A general issue for both reductions that arises when moving from the unstructured to the module setting is that we have two different possibilities to define what a *binary secret* exactly means. Recall that a secret underlying an M-LWE instance is a vector over the dual R^\vee . Thus, when interpreting it as a vector over the integers \mathbb{Z} to define a binary notion, we can either use the coefficient embedding τ or the canonical embedding σ , both introduced in Section 1.1. We argue in Section 2.5 for the case of power-of-2 cyclotomics, that using the canonical embedding

for binary secrets requires the rank d to be larger by a factor n than when using the coefficient embedding. This makes the definition via the canonical embedding impractical and we thus use the coefficient embedding to define a binary secret. Additionally, the coefficient embedding is also the favorite choice once a bin-M-LWE-based scheme is used in practice and thus needs to be implemented. When manipulating Gaussians for theoretical results, however, the canonical embedding is much more suitable. This requires to switch between both embeddings, which in turn has an impact on the parameters overall.

Table 2.1: Comparison between both reductions, where n denotes the ring degree, q the modulus, m the number of samples and d the module rank

Property	Simple (Section 2.2)	Improved (Section 2.3)
LWE analogue	[GKPV10] using RD	[BLP ⁺ 13]
minimal rank d	$\log_2 q + \Omega(\log_2 n)$	$2 \log_2 q + \omega(\log_2 n)$
noise ratio	$O(\sqrt{mn^2 d})$	$O(n^2 \sqrt{d})$
class of number fields	cyclotomics	cyclotomics
sample dependency	dependent	independent
modulus preserving	yes	yes
ring degree preserving	yes	yes
conditions on modulus q	prime	number-theoretic restrictions
decision/search variant	search	decision
complexity of proof	simple	involved

The main challenge of the improved reduction from Section 2.3 is the use of matrices over the ring of integers R , as opposed to matrices over \mathbb{Z} for LWE. The proof in [BLP⁺13, Lem. 4.7] requires the construction of unimodular matrices, which is not straightforward to adapt in the module setting because units of the quotient ring $R_q = R/qR$ are much harder to describe than the units of $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$. This is the reason why we need to control the splitting structure of the cyclotomic polynomial modulo q . Lemma 1.2 [LS18, Thm. 1.1] solves this issue but 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 .

Both reductions are modulus and ring degree preserving, but so far limited to cyclotomic fields. This restriction stems from the difficulty to capture the invertibility of elements and matrices in more general number fields (needed in Lemma 2.1 and Lemma 2.2) and to bound the norm of the Vandermonde matrix when switching from the coefficient to the canonical embedding (used in Theorem 2.1). Another weaker restriction comes from the LHL (Lemma 1.23) which is used in both reductions and only shown for number fields $K = \mathbb{Q}(\zeta)$ such that the modulus q is coprime with the index $[R : \mathbb{Z}[\zeta]]$, where R is the associated ring of integers of K . This class contains all cyclotomic number fields as in this case $R = \mathbb{Z}[\zeta]$. All other results are proven for general number fields.

The simple reduction from Section 2.2 uses the Rényi divergence (RD) as a measure of distance of two probability distributions, which requires to fix the number of samples beforehand and restricts the results to the corresponding search versions. As the improved reduction only uses the statistical distance, and not the Rényi divergence, both constraints are not needed there.

As an additional contribution, which did not appear in the two original publications [BJRW20, BJRW21], we show that both reductions generalize to larger secret distributions. More precisely, we show in Section 2.4 that the problem η -M-LWE, a variant of M-LWE where the secret's coefficients are sampled uniformly at random over R_η^\vee with $R_\eta = \{x \in R : \tau(x) \in \{0, \dots, \eta - 1\}^n\}$, is at least as hard as standard M-LWE. The problem bin-M-LWE is the special case of η -M-LWE

with $\eta = 2$. In both reductions, a larger η leads to a weaker requirement on the module rank d , allowing smaller ranks by a factor of roughly $1/\log_2 \eta$. However, a larger η increases the noise ratio by a factor $(\eta-1)/\log_2 \eta$ in the warm-up reduction from Section 2.2 and by a factor $(\eta-1)^2/(\sqrt{\log_2 \eta})$ in the improved reduction from Section 2.3.

2.1.2 Related Work

Extended M-LWE

In Section 2.3, we introduce an intermediate problem, which we call extended M-LWE. In this variant, for an instance of M-LWE given by $(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$, additional information on the noise vector \mathbf{e} in form of an inner-product $\langle \mathbf{e}, \mathbf{z} \rangle$ with some known hint vector \mathbf{z} is given. In the work of Alperin-Sheriff and Apon [AA16] for their reduction from M-LWE to the deterministic variant Module Learning With Rounding, a variant of M-LWE that is also called extended M-LWE is introduced. In contrast to the full inner product as in our definition, the extended version in [AA16] gives the *trace* of the inner product $\text{Tr}(\langle \mathbf{e}, \mathbf{z} \rangle)$ to the attacker. This variant is not suited for our reduction as the field trace does not provide enough information for our lossy argument in Lemma 2.5. We discuss further differences in Section 2.3.1 and Section 2.3.2. A very recent work by Lyubashevsky et al. [LNS21] defines yet another variant of M-LWE in the context of lattice-based zero-knowledge proofs that is called extended M-LWE. Instead of the full inner product as in our definition, they provide only the *sign* of the inner product as an additional hint for the attacker. Again, this is not sufficient for our lossy argument in Lemma 2.5.

Entropic Secrets

Another line of research which studies secret distributions that are not uniform over \mathbb{Z}_q^n , is the one on LWE using so-called *entropic* secrets [GKPV10, AKPW13, BD20a]. This variant encompasses not only secrets within a ball of small radius, but also more general secret distributions other than the uniform distribution, studying it from an entropic point of view. This allows to take into account possible leakage on the secret, even though it may come from a distribution yielding vectors of large norms. Recently, the entropic secret setting has also been investigated for R-LWE [BD20b] and M-LWE [LWW20]. As an independent contribution, Lin et al. [LWW20] provide an alternative LHL over rings. The LHL we use in our proof (Lemma 1.23) with respect to the statistical distance is an adaptation of [Mic07, Thm. 4.2] to the ring setting and requires the rank to be increased by a super-logarithmic factor $\log_2 q + \omega(\log_2 n)$. In the LHL by Lin et al. [LWW20, Thm. 1], the rank only needs to be larger by a logarithmic factor $2 \log_2 q$. However, they require the defining polynomial of the underlying number field to be irreducible modulo q , which adds strong restrictions on the modulus q . For example, for the ν -th cyclotomic number field, where ν is a power of 2, the defining polynomial is reducible for every positive integer q and thus there is no such q that fulfills the restrictions.

Binary Noise

Instead of choosing a binary secret for LWE, one can also look at the variant of LWE, where the *noise* is sampled over the set of binary vectors. For instance, this problem is studied in the context of efficient encryption schemes for constrained devices. It was introduced by Micciancio and Peikert [MP13] together with a reduction from worst-case lattice problems as long as the number of samples is only slightly larger than the LWE dimension n . Further work is needed to investigate if this reduction can be adapted to structured variants. There is a sequence of works showing that binary error LWE (and more generally LWE with small errors)

becomes easy to solve via lattice algorithms or algebraic algorithms if sufficiently many samples are available [AG11, MP13, ACF⁺15, KF15, BGPW16, STA20]. More concretely, there is a polynomial-time algorithm solving LWE with a binary noise if the number of provided samples is quadratic in the dimension n [AG11]. In this thesis, we are not studying this variant.

Practical Hardness

In this chapter, we are studying the *theoretical* hardness of M-LWE with binary secrets by providing a reduction from M-LWE with uniform secrets, which thus connects it via [LS15] to worst-case problems over module lattices. As often the case for lattice-based cryptography, there is a gap between the theoretical hardness (supported by reductions as in this chapter) and the *practical* hardness of the problem. By practical hardness we mean the concrete difficulty to solve the problem using the best known algorithms.

As the concrete difficulty also serves to select sample parameters for schemes that rely on the hardness of LWE, there are handy tools that allow to estimate the practical hardness of LWE, depending on its different parameters. Using for instance the LWE Estimator by Albrecht et al. [APS15] publicly available at <https://bitbucket.org/malb/lwe-estimator/src>, one can measure the practical hardness of LWE with a binary secret. As there are no algorithms that exploit the structure of M-LWE, the following reasoning holds for bin-M-LWE as well.

More concretely, using the computation model of Alkim et al. [ADPS16] with respect to the core SVP model, we see in Figure 2.1 that in order to keep the same estimated security level for bin-LWE as for LWE, it suffices to increase the LWE dimension n and the Gaussian width α by a multiplicative factor even less than 2. Here, we use the parameters recommended by Regev’s reduction [Reg09] for $n = 1024$, i.e., the modulus $q = n^2$ and the Gaussian width $\alpha = 1/(\sqrt{n} \log_2(n)^2)$. In the first estimates we set `secret_distribution = uniform` and in the second estimates we set `secret_distribution = (0, 1)`, which sets it to randomly chosen binary secrets. We see that in the first run, the best algorithm is the one that solves the associated u-SVP problem, guaranteeing 379 bits of security. In the second run, the best algorithm is again the one that solves the associated u-SVP problem, guaranteeing even 492 bits of security. The only difference to both executions of the estimator is that we increased n by a factor [1.8] and α by a factor 1.8. This is much smaller than the (super-)logarithmic factor for the dimension and the \sqrt{n} factor for the Gaussian width that Brakerski et al. [BLP⁺13] obtain for LWE and that we obtain for M-LWE.

```
sage: load("estimator.py")
sage: n,alpha,q=Param_Regev(1024)
sage: _ = estimate_lwe(n,alpha,q,secret_distribution=uniform,reduction_cost_model=BKZ_ADPS16)
usvp: rop: ≈24379.3, red: ≈24379.3, s_0: 1.001672, β: 1299, d: 2851, m: ≈2411.5
dec: rop: ≈24516.6, m: ≈2411.6, red: ≈24516.6, s_0: 1.001367, β: 1683, d: 3196, baba1: ≈24504.3, baba1_op: ≈24519.4, repeat: ≈2425.2, ε: ≈24-23.0
dual: rop: ≈24525.7, m: ≈2411.6, red: ≈24525.7, s_0: 1.001443, β: 1571, d: 3138, |v|: ≈2413.1, repeat: ≈2467.0, ε: ≈24-32.0
sage: _ = estimate_lwe(round(1.8)*n,1.8*alpha,q,secret_distribution=(0,1),reduction_cost_model=BKZ_ADPS16)
Warning: the LWE secret is assumed to have Hamming weight 1024.
usvp: rop: ≈24492.0, red: ≈24492.0, s_0: 1.001366, β: 1685, d: 3371, m: 1322, repeat: 1, k: 0, postprocess: 0
dec: rop: ≈241453.7, m: ≈2412.1, red: ≈241453.7, s_0: 1.000604, β: 4642, d: 6302, baba1: ≈241111.8, baba1_op: ≈241nf, repeat: ≈2498.2, ε: ≈24-96.0
dual: rop: ≈24547.8, m: 1508, red: ≈24547.8, s_0: 1.001255, β: 1876, repeat: ≈24473.0, d: 3492, c: 470.615, k: 64, postprocess: 1
sage: _
```

Figure 2.1: Screenshot of running the LWE Estimator [APS15] with SageMath on my laptop.

Kirchner and Fouque [KF15] improve the so-called BKW algorithm to solve the LWE problem and, in particular, run some experiments on the binary secret variant, assuming that enough samples are provided. Chen et al. [CCLS20] systematically run experiments on binary secret LWE with dimensions at most 200 to better understand the concrete behavior of lattice algorithms such as BKZ for this special variant. However, their range of parameters does not cover the ones used by cryptographic constructions (where the LWE dimension is usually starting from 500).

2.1.3 Roadmap

The chapter is structured as follows. In Section 2.2, we prove the first reduction for the hardness of M-LWE with a binary secret, which we call the warm-up reduction. In Section 2.3, we provide an improved reduction to show the hardness of M-LWE with a binary secret, which needs two new intermediate problems, called first-is-errorless M-LWE and ext-M-LWE. Furthermore, we generalize both reductions in Section 2.4 to secrets of bounded norm, not necessarily binary secrets. We conclude the chapter by arguing in Section 2.5 that the coefficient embedding is the suited choice in the setting of binary secrets.

2.2 Warm-up: A Simple Reduction

In the following, we show a first reduction for the hardness of M-LWE with a binary secret. Our proof follows the original proof structure of Goldwasser et al. [GKPV10] in the case of LWE, while achieving much better noise parameters by using the Rényi divergence instead of the statistical distance. The improvement on the noise rate 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). A similar effect arises with respect to the rank condition in comparison with the second proof for bin-M-LWE that we prove in the next section. More precisely, as we use the LHL with respect to the Rényi divergence, and not the statistical distance, we can have a rank that is logarithmic in the ring degree n , instead of super-logarithmic. Note that we only recently found out that using the LHL with respect to the Rényi divergence decreases the rank condition from super-logarithmic down to logarithmic.² This is why this didn't appear in the original publication [BJRW20, Thm. 1] and can be seen as an improvement achieved afterwards. However, using the Rényi divergence as a tool for distance measurement requires to move to the *search* variants of M-LWE (denoted by M-SLWE) and its binary version, respectively. Additionally, it asks to fix the number of samples a priori.

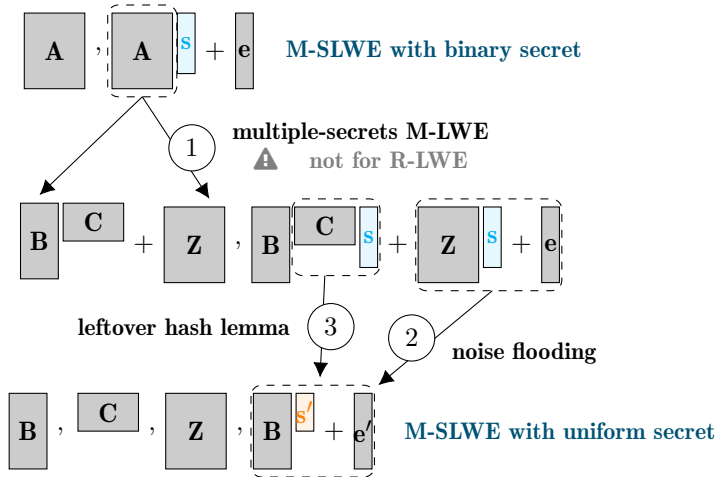


Figure 2.2: Overview of the proof of Theorem 2.1.

²We would like to thank Thomas Prest for pointing out to us the connection between the collision probability and the Rényi divergence.

To facilitate the understanding of this section, we start by illustrating the high level idea of the proof in Figure 2.2 (where the Latex/tikz credits go to my co-author Corentin Jeudy). Given an instance of bin-M-LWE by $(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$, our goal is to transform it to a related instance of M-LWE defined by $(\mathbf{B}, \mathbf{B} \cdot \mathbf{s}' + \mathbf{e}')$. Note that the secret \mathbf{s} is modulo 2 and the secret \mathbf{s}' is modulo q . At the core of the hardness proof of bin-M-LWE lies a lossy argument, where the public matrix \mathbf{A} is replaced by a lossy matrix $\mathbf{B} \cdot \mathbf{C} + \mathbf{Z}$, which corresponds to the second part of some multiple-secrets M-LWE sample. Note that the rank of the matrix \mathbf{B} is smaller than the one of the matrix \mathbf{A} , motivating the description *lossy*. Here, we can see that this argument does not work for R-LWE (which corresponds to M-LWE of rank 1) as there is no possibility to replace the public matrix consisting of one column by a matrix of smaller rank. To argue that an adversary cannot distinguish between the two cases, we need the hardness of the *decision* M-LWE problem as well. In a second step, the term $\mathbf{Z} \cdot \mathbf{s} + \mathbf{e}$ is replaced by the new noise \mathbf{e}' , where the Rényi divergence of both expressions can be bounded by a constant using the properties of the Rényi divergences of Gaussian distributions. This step is commonly called noise flooding. Finally in the third step, the product $\mathbf{C} \cdot \mathbf{s}$ is replaced by the uniform secret \mathbf{s}' , where the Rényi divergence of both elements can be bounded by a constant using a variant of the Leftover Hash Lemma (LHL). The use of the LHL is also the reason why our reduction only works for module ranks larger than $\log_2(q) + \Omega(\log_2 n)$. Informally speaking, it requires the ratio between the number of rows of \mathbf{C} and its number of columns to be logarithmic in order to bound the Rényi divergence. As providing the full information of \mathbf{B}, \mathbf{C} and \mathbf{Z} only makes the problem easier, we end up with some standard M-LWE instance.

We now formally prove the hardness of M-SLWE with a binary secret. Overall, our results are restricted to cyclotomic number fields. As in Section 1.1.5, we denote by λ the scaling factor to map elements from R^\vee to R . We study the M-LWE problem in its discrete version, as presented in Section 1.4.3. The theorem uses the discrete Gaussian distribution $\psi = \mathcal{D}_{R^\vee, \alpha q}$ for some positive real α and the smoothing parameter η_ε , as defined in Section 1.3.1.

Theorem 2.1 (Binary Hardness of M-LWE - Warm-up)

Let K be a cyclotomic number field of degree n with R its ring of integers. Let $k, d, m, q \in \mathbb{N}$ such that q is prime and m polynomial in n . Further, let α and β be positive real numbers such that $\beta \geq \sqrt{m} \cdot n^2 d \cdot \alpha$. Let $\varepsilon \in \mathbb{R} \cap [0, 1/2)$ such that $\beta q \geq \eta_\varepsilon(R^\vee)$ and $\varepsilon = O(\frac{1}{m})$. Then, for any $d \geq k \cdot \log_2 q + \Omega(\log_2 n)$, there is a probabilistic polynomial-time reduction from $\text{M-SLWE}_{n,k,q,\mathcal{D}_{R^\vee,\beta q}}^m$ and $\text{M-LWE}_{n,k,q,\mathcal{D}_{R^\vee,\alpha q}}^{m,d}$ to $\text{bin-M-SLWE}_{n,d,q,\mathcal{D}_{R^\vee,\beta q}}^m$.

The degree n of the number field K , the modulus q and the number of samples m are preserved. The reduction increases the rank of the module from k to $k \cdot \log_2 q + \Omega(\log_2 n)$ and the Gaussian width from αq to $\alpha q \cdot \sqrt{m} \cdot n^2 d$. Further, $\text{M-LWE}_{n,k,q,\mathcal{D}_{R^\vee,\alpha q}}^{m,d}$ trivially reduces to $\text{M-SLWE}_{n,k,q,\mathcal{D}_{R^\vee,\beta q}}^m$, as $\beta \geq \alpha$.

Proof: Given a $\text{bin-M-SLWE}_{n,d,q,\mathcal{D}_{R^\vee,\beta q}}^m$ sample $(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}) \in (R_q)^{m \times d} \times (R_q^\vee)^m$, with $\mathbf{s} \in (R_2^\vee)^d$ and $\mathbf{e} \leftarrow (\mathcal{D}_{R^\vee,\beta q})^m$, the search problem asks to find \mathbf{s} and \mathbf{e} . In order to prove the statement, we define different hybrid distributions:

- $H_0 : (\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e})$, as in $\text{bin-M-SLWE}_{n,d,q,\mathcal{D}_{R^\vee,\beta q}}^m$,
- $H_1 : (\mathbf{A}' = \lambda(\mathbf{B}\mathbf{C} + \mathbf{Z}), \mathbf{A}' \cdot \mathbf{s} + \mathbf{e})$, where $\mathbf{B} \leftarrow U((R_q)^{m \times k})$, $\mathbf{C} \leftarrow U((R_q^\vee)^{k \times d})$, and $\mathbf{Z} \leftarrow$

$(\mathcal{D}_{R^\vee, \alpha q})^{m \times d}$ and \mathbf{s}, \mathbf{e} as in H_0 ,

- $H_2 : (\mathbf{B}, \mathbf{C}, \mathbf{Z}, \mathbf{B}(\lambda \mathbf{C}\mathbf{s}) + \mathbf{Z}(\lambda \mathbf{s}) + \mathbf{e})$, where $\mathbf{B}, \mathbf{C}, \mathbf{Z}, \mathbf{s}, \mathbf{e}$ as in H_1 ,
- $H_3 : (\mathbf{B}, \mathbf{C}, \mathbf{Z}, \mathbf{B}(\lambda \mathbf{C}\mathbf{s}) + \mathbf{e}')$, where $\mathbf{e}' \leftarrow (\mathcal{D}_{R^\vee, \beta q})^m$ and $\mathbf{B}, \mathbf{C}, \mathbf{Z}, \mathbf{s}$ as in H_2 ,
- $H_4 : (\mathbf{B}, \mathbf{C}, \mathbf{Z}, \mathbf{B}\mathbf{s}' + \mathbf{e}')$, where $\mathbf{s}' \leftarrow U((R^\vee)^k)$ and $\mathbf{B}, \mathbf{C}, \mathbf{Z}, \mathbf{e}'$ as in H_3 .

For $j \in \{0, \dots, 4\}$, we denote by P_j the problem of finding the secret \mathbf{s} (resp. \mathbf{s}' in H_4), given a sample of the distribution H_j . We say that problem P_j is hard if for any probabilistic polynomial-time attacker \mathcal{A} the advantage of solving P_j is negligible, thus $\text{Adv}_{P_j}[\mathcal{A}(H_j) = \mathbf{s}] \leq n^{-\omega(1)}$, where n is the degree of K . The overall idea is to show that if P_4 is hard, then P_0 is hard as well.

Problem P_4 is hard. By the hardness assumption of M-SLWE $_{n,k,q,\mathcal{D}_{R^\vee, \beta q}}^m$, it yields

$$\text{Adv}_{P_4}[\mathcal{A}(H_4) = \mathbf{s}'] \leq n^{-\omega(1)}.$$

From P_4 to P_3 . By the probability preservation property of the Rényi divergence (Lemma 1.20), we have

$$\text{Adv}_{P_3}[\mathcal{A}(H_3) = \mathbf{s}]^2 \leq \text{Adv}_{P_4}[\mathcal{A}(H_4) = \mathbf{s}'] \cdot \text{RD}_2(H_3 \| H_4).$$

The only difference between the distributions H_3 and H_4 is that the element $\lambda \mathbf{C}\mathbf{s}$ in H_3 is replaced by a uniform $\mathbf{s}' \in (R^\vee)^k$ in H_4 . Our aim is to show that their Rényi divergence can be bounded above by a constant. We set $\tilde{\mathbf{C}} = \lambda \mathbf{C} \in (R_q)^{k \times d}$ and $\tilde{\mathbf{s}} = \lambda \mathbf{s} \in (R_2)^d$. By the Ring Leftover Hash Lemma stated in Lemma 1.23 with $\eta = 2$, the Rényi divergence between the distribution $(\tilde{\mathbf{C}}, \tilde{\mathbf{s}})$ and the distribution $(\tilde{\mathbf{C}}, \tilde{\mathbf{s}}')$ is bounded above by $(1 + q^k/2^d)^n$. Note that for cyclotomic fields, it yields $R = \mathbb{Z}[\zeta]$ and thus the index $[R : \mathbb{Z}[\zeta]] = 1$ is always coprime to q . Dividing the first and the second part of both distributions by λ preserves the Rényi divergence. As we require $d \geq k \log_2 q + \Omega(\log_2 n)$, we obtain $\text{RD}_2(H_3 \| H_4) \leq (1 + O(1/n))^n$, which can be asymptotically bounded by a constant.

From P_3 to P_2 . Again, by the probability preservation property of the Rényi divergence (Lemma 1.20), we obtain

$$\text{Adv}_{P_2}[\mathcal{A}(H_2) = \mathbf{s}]^2 \leq \text{Adv}_{P_3}[\mathcal{A}(H_3) = \mathbf{s}] \cdot \text{RD}_2(H_2 \| H_3).$$

In order to compute the Rényi divergence of H_2 and H_3 , we need to compute the Rényi divergence of $\mathbf{Z}(\lambda \mathbf{s}) + \mathbf{e}$ and \mathbf{e}' . Using the noise flooding bound of Lemma 1.14 with $\eta = 2$, we know that each of the m coefficients of $\mathbf{Z}(\lambda \mathbf{s})$ is bounded above by $\alpha q d n^2$ with probability $1 - 2^{-\Omega(n)}$. Thus, it suffices to compute the Rényi divergence of $(\mathcal{D}_{R^\vee, \beta q, c})^m$ and $(\mathcal{D}_{R^\vee, \beta q})^m$, where $c \in R^\vee$ with norm bounded above by $\alpha q d n^2$. Using that $\beta q \geq \eta_\varepsilon(R^\vee)$, the multiplicativity of the Rényi divergence (Lemma 1.20) and the result of Lemma 1.21 about the Rényi divergence of

shifted discrete Gaussians, we deduce

$$\begin{aligned} \text{RD}_2((\mathcal{D}_{R^\vee, \beta q, c})^m \| (\mathcal{D}_{R^\vee, \beta q})^m) &= \text{RD}_2(\mathcal{D}_{R^\vee, \beta q, c} \| \mathcal{D}_{R^\vee, \beta q})^m \\ &\leq \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^{2m} \cdot \exp\left(\frac{2\pi\|c\|^2}{(\beta q)^2}\right)^m \\ &\leq \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^{2m} \cdot \exp(2\pi). \end{aligned}$$

The last inequality comes from the restriction $\frac{\alpha}{\beta} \leq \frac{1}{\sqrt{m \cdot n^2 d}}$ which leads to $\frac{2\pi\|c\|^2}{(\beta q)^2} \leq \frac{2\pi}{m}$. Here, we simply use that the function \exp is increasing for increasing input.

For the Rényi divergence to be bounded by a constant, we also need $\varepsilon = O(\frac{1}{m})$. Indeed, we have $\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^2 = \left(1 + \frac{4\varepsilon/(1-\varepsilon)}{2}\right)^2 < \exp\left(\frac{4\varepsilon}{1-\varepsilon}\right)$ as $\left(1 + \frac{x}{y}\right)^y < \exp(x)$ for any $x, y > 0$. As we require $\varepsilon < \frac{1}{2}$, it yields $\frac{1}{1-\varepsilon} < 2$ and thus, we get $\left(\frac{1+\varepsilon}{1-\varepsilon}\right)^{2m} < \exp(8m\varepsilon)$ and therefore $\varepsilon = O(\frac{1}{m})$ suffices.

From P_2 to P_1 . Since more information is given in distribution H_2 than in distribution H_1 , the problem P_1 is harder than P_2 and hence

$$\text{Adv}_{P_1}[\mathcal{A}(H_1) = \mathbf{s}] \leq \text{Adv}_{P_2}[\mathcal{A}(H_2) = \mathbf{s}].$$

From P_1 to P_0 . By the hardness assumption of $\text{M-LWE}_{n,k,q,\mathcal{D}_{R^\vee, \alpha q}}^{m,d}$, the distributions H_0 and H_1 are computationally indistinguishable. More concretely,

$$\begin{aligned} \text{Adv}_{P_0}[\mathcal{A}(H_0) = \mathbf{s}] &\leq \text{Adv}_{P_1}[\mathcal{A}(H_1) = \mathbf{s}] + d \cdot \text{Adv}_{\text{M-LWE}} \\ &\leq \text{Adv}_{P_1}[\mathcal{A}(H_1) = \mathbf{s}] + d \cdot n^{-\omega(1)}, \end{aligned}$$

where d is the number of secret vectors, represented as the columns of the matrix \mathbf{C} . Putting all equations from above together, we obtain

$$\begin{aligned} \text{Adv}_{P_0}[\mathcal{A}(H_0) = \mathbf{s}] &\leq \text{Adv}_{P_1}[\mathcal{A}(H_1) = \mathbf{s}] + d \cdot \text{Adv}_{\text{M-LWE}} \\ &\leq \text{Adv}_{P_2}[\mathcal{A}(H_2) = \mathbf{s}] + d \cdot \text{Adv}_{\text{M-LWE}} \\ &\leq \sqrt{\text{Adv}_{P_3}[\mathcal{A}(H_3) = \mathbf{s}] \cdot \text{RD}_2(H_2 \| H_3)} + d \cdot \text{Adv}_{\text{M-LWE}} \\ &\leq \sqrt{\sqrt{\text{Adv}_{P_4}[\mathcal{A}(H_4) = \mathbf{s}']} \cdot \text{RD}_2(H_3 \| H_4)} \cdot \text{RD}_2(H_2 \| H_3) \\ &\quad + d \cdot \text{Adv}_{\text{M-LWE}} \\ &\leq n^{-\omega(1)}. \end{aligned}$$

As the problem P_0 is exactly $\text{bin-M-SLWE}_{n,d,q,\mathcal{D}_{R^\vee, \beta q}}^m$, this concludes the proof. ■

2.3 An Improved Reduction

In this section, we improve the hardness result of the section above for M-LWE with binary secrets for cyclotomic fields. Our proof follows the same idea as in [BLP⁺13] that we adapt over modules. More precisely, we show a reduction from M-LWE with rank k to bin-M-LWE

with rank d satisfying $d \geq (k+1)\log_2 q + \omega(\log_2 n)$. The reduction preserves the modulus q , that needs to be prime satisfying number-theoretic restrictions (as detailed in Theorem 2.2), the ring degree n and the number of samples m , but the noise is increased by a factor of $n\sqrt{2d}\sqrt{4n^2+1}$. The noise ratio is polynomial in n , but smaller than the ratio $n^2 d \sqrt{m}$ that we obtained with the first proof from Section 2.2. Thus, we improve the noise parameters by a factor of \sqrt{dm} , which is advantageous as the typical choice for m in theoretical reductions is $m = \Theta(n \log_2 n)$. For the reduction, m also needs to be larger than the target module rank d , and at most polynomial in n because of the hybrid argument used in Lemma 2.4. As the reduction does not need the Rényi divergence as a measure for distance, the results hold for the *decision* versions of M-LWE and bin-M-LWE, as opposed to the reduction from the section above. On the other hand, the improved noise parameters come at the cost of a larger rank condition (super-logarithmic instead of logarithmic) and a more complex proof, involving several intermediate problems. Theorem 2.2 holds for all cyclotomic fields, whereas most results even apply for all number fields $K = \mathbb{Q}(\zeta)$ such that the ring of integers is $R = \mathbb{Z}[\zeta]$, the bottleneck being the construction in Lemma 2.2.

We now state our main result of this section, first for general cyclotomic number fields and then specialized to power-of-2 cyclotomics. We illustrate the sequence of reductions needed to prove Theorem 2.2 and Corollary 2.1 in Figure 2.3. The formal definitions of the extended and first-is-errorless variants of M-LWE are given in Section 2.3.1 and Section 2.3.2 below. We refer to Section 1.3.2 for the formal definitions of the different Gaussian distributions over a number field K that arise in this section, as for instance D_α , $\mathcal{D}_{R,\alpha}$ or $\Psi_{\leq\alpha}$.

Theorem 2.2 (Binary Hardness of M-LWE)

Let $\nu \in \mathbb{N}$ with prime-power factorization $\nu = \prod_j p_j^{e_j}$. Further, let K be the ν -th cyclotomic field of degree $n = \varphi(\nu)$, and R its ring of integers. Set $\mu = \prod_j p_j$ and let q be a prime number such that $q \equiv 1 \pmod{\mu}$, $\text{ord}_\nu(q) = \nu/\mu$ and $q > \max(2n, \mathfrak{s}_1(\mu)^{\varphi(\mu)})$, where $\mathfrak{s}_1(\mu)$ denotes the largest singular value of the Vandermonde matrix of the μ -th cyclotomic field. Further, let $k, d, m \in \mathbb{N}$ such that $d \geq (k+1)\log_2 q + \omega(\log_2 n)$, and $d \leq m \leq \text{poly}(n)$. Let $\alpha \geq q^{-1}\sqrt{\ln(2nd(1+1/\varepsilon))/\pi}$ and $\beta \geq \alpha \cdot n\sqrt{2d}\sqrt{4n^2+1}$. Then there is a reduction from $\text{M-LWE}_{n,k,q,D_\alpha}^m$ to $\text{bin-M-LWE}_{n,d,q,\Psi_{\leq\beta}}^m$, such that if \mathcal{A} solves the latter with advantage $\text{Adv}[\mathcal{A}]$, then there exists an algorithm \mathcal{B} that solves the former with advantage

$$\text{Adv}[\mathcal{B}] \geq \frac{1}{3m} \left(\text{Adv}[\mathcal{A}] - \frac{1}{2} \sqrt{\left(1 + \frac{q^{k+1}}{2^d}\right)^n - 1} \right) - \frac{37\varepsilon}{2}.$$

As mentioned before, the noise ratio β/α is given by the term $n\sqrt{2d}\sqrt{4n^2+1}$, which is composed of three different parts. The factor n encapsulates the norm distortion between the coefficient embedding τ and the canonical embedding σ , as well as the actual length of the binary vectors. The second term $\sqrt{2d}$ stems from the masking of \mathbf{z} when introduced in the first hybrid in the proof of Lemma 2.5. The last factor $\sqrt{4n^2+1}$ solely represents the impact of giving information on the error in the ext-M-LWE problem.

Corollary 2.1 (Binary Hardness of M-LWE for Power-of-2 Cyclotomics)

Let $\nu \in \mathbb{N}$ be a power of 2 and K be the ν -th cyclotomic field of degree $n = \nu/2$, and R its ring of integers. Let $q \in \mathbb{N}$ be prime such that $q \equiv 5 \pmod{8}$. Further, let $k, d, m \in \mathbb{N}$ such

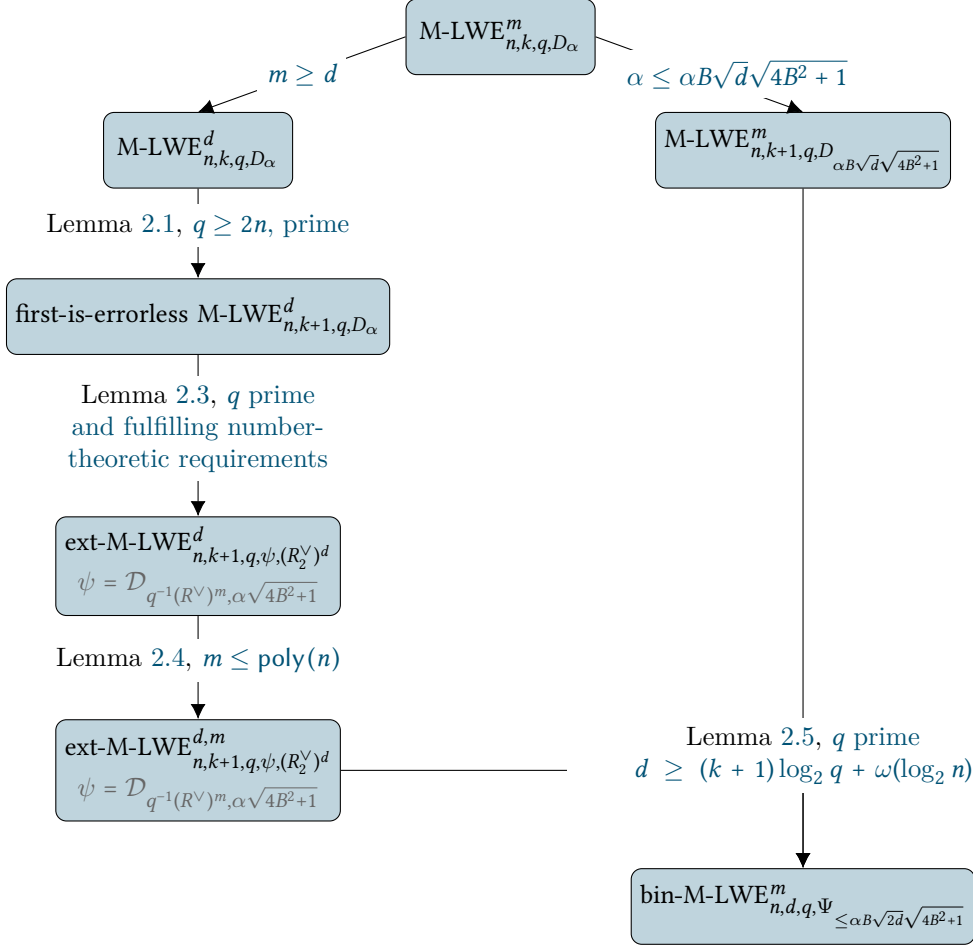


Figure 2.3: Overview of the proof of Theorem 2.2, starting from M-LWE with uniform secrets via first-is-errorless M-LWE and ext-M-LWE to M-LWE with binary secrets. The bound B is defined as $B = \max_{x \in R_2} \|\sigma(x)\|_\infty$, where σ is the canonical embedding. In cyclotomic fields, we have $B \leq n$. Note that Lemma 2.5 uses d samples from ext-M-LWE, where d is the module rank in bin-M-LWE. The number-theoretic assumptions on q concern the splitting behavior of the cyclotomic polynomial in $\mathbb{Z}_q[x]$, and are discussed in Section 2.3.2.

that $d \geq (k+1) \log_2 q + \omega(\log_2 n)$, and $d \leq m \leq \text{poly}(n)$. Let $\alpha \geq q^{-1} \sqrt{\ln(2nd(1+1/\varepsilon))/\pi}$ and $\beta \geq \alpha \cdot n \sqrt{2d} \sqrt{4n^2 + 1}$. Then there is a reduction from $\text{M-LWE}_{n,k,q,D_\alpha}^m$ to $\text{bin-M-LWE}_{n,d,q,\Psi \leq \beta}^m$, such that if \mathcal{A} solves the latter with advantage $\text{Adv}[\mathcal{A}]$, then there exists an algorithm \mathcal{B} that solves the former with advantage

$$\text{Adv}[\mathcal{B}] \geq \frac{1}{3m} \left(\text{Adv}[\mathcal{A}] - \frac{1}{2} \sqrt{\left(1 + \frac{q^{k+1}}{2^d}\right)^n - 1} \right) - \frac{37\varepsilon}{2}.$$

2.3.1 First-is-errorless M-LWE

We follow the same idea as Brakerski et al. [BLP⁺13] by gradually giving more information to the adversary while proving that this additional information does not increase the advantage too much. We define the module version of first-is-errorless LWE, from [BLP⁺13], where the first equation is given without error. A similar definition and reduction from M-LWE are given in [AA16]. The only difference between the two reductions comes from the pre-processing step, which is simplified in our case due to the further restrictions on q of our overall reduction.

For a number field K of degree n and with ring of integers R , we set $R_q = R/qR$, $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$, and $\mathbb{T}_{R^\vee} = K_{\mathbb{R}}/R^\vee$ as usual. Further, recall the definition of the M-LWE distribution $A_{\mathbf{s},\psi}^{\mathcal{M}}$ from Section 1.4.3.

Definition 2.1 (First-is-errorless M-LWE): Let $n, k, q \in \mathbb{N}$ and K be a number field of degree n and R its ring of integers. Let Υ be a distribution over a family of distributions over $K_{\mathbb{R}}$. The first-is-errorless M-LWE $_{n,k,q,\Upsilon}$ problem is to distinguish between the following cases. On the one hand, the first sample is uniform over $(R_q)^k \times q^{-1}R^\vee/R^\vee$ and the rest are uniform over $(R_q)^k \times \mathbb{T}_{R^\vee}$. On the other hand, there is some unknown $\mathbf{s} \leftarrow U((R_q^\vee)^k)$ and $\psi \leftarrow \Upsilon$ such that the first sample is from $A_{\mathbf{s},\{0\}}^{\mathcal{M}}$ and the rest are distributed as $A_{\mathbf{s},\psi}^{\mathcal{M}}$, where $\{0\}$ is the distribution that is deterministically 0.

If the number $m \in \mathbb{N}$ of samples is fixed, we write first-is-errorless M-LWE $_{n,k,q,\Upsilon}^m$. The following lemma shows for cyclotomic fields that first-is-errorless M-LWE with rank k is at least as hard as the standard M-LWE problem with rank $k-1$.

Lemma 2.1 (M-LWE to first-is-errorless M-LWE, adapted from [BLP⁺13, Lem. 4.3])

Let K be the ν -th cyclotomic field of degree $n = \varphi(\nu)$, and R its ring of integers. Let $q \geq 2n$ be a prime integer such that $q \nmid \nu$, k a positive integer, and Υ a distribution over a family of distributions over $K_{\mathbb{R}}$. There is PPT reduction from M-LWE $_{n,k-1,q,\Upsilon}$ to first-is-errorless M-LWE $_{n,k,q,\Upsilon}$.

Proof: *Pre-processing:* The reduction first chooses $\mathbf{a}' \leftarrow U((R_q)^k)$ and then $\mathbf{b}_2, \dots, \mathbf{b}_k$ independently from $U((R_q)^k)$ such that $\mathbf{a}', \mathbf{b}_2, \dots, \mathbf{b}_k$ are R_q -linearly independent. Each time we draw a uniformly random column, the probability that the new column is R_q -linearly independent with the previous ones is at least $1 - n/q$ for $q \geq n$ by Lemma 1.3. Since we

require $q \geq 2n$, this probability is at least $1/2$. Therefore, we only need a polynomial number of uniformly sampled columns in R_q^k to construct a matrix of $R_q^{k \times k}$ invertible modulo qR . The preprocessing step results in a matrix $\mathbf{U} = [\mathbf{a}', \mathbf{b}_2, \dots, \mathbf{b}_k] \in (R_q)^{k \times k}$ that is invertible modulo qR according to Lemma 1.3.

Reduction: Then, sample s_0 uniformly in R_q^\vee . The reduction is as follows. For the first sample, it outputs $(\mathbf{a}', q^{-1} \cdot s_0 \bmod R^\vee) \in (R_q)^k \times q^{-1}R^\vee/R^\vee$. The other samples are produced by taking $(\mathbf{a}, b) \in (R_q)^{k-1} \times \mathbb{T}_{R^\vee}$ from the M-LWE challenger, picking a fresh randomly chosen $a'' \in R_q$, and outputting $(\mathbf{U}(a''|\mathbf{a}), b + q^{-1}(s_0 \cdot a'')) \bmod R^\vee \in (R_q)^k \times \mathbb{T}_{R^\vee}$, with the vertical bar denoting concatenation. We now analyze correctness.

First note that the first component is uniform over $(R_q)^k$. Indeed, \mathbf{a}' is uniform over $(R_q)^k$ for the first sample, and since \mathbf{a} is uniform over $(R_q)^{k-1}$, a'' is uniform over R_q , and \mathbf{U} is invertible in $(R_q)^{k \times k}$, then $\mathbf{U}(a''|\mathbf{a})$ is uniform over $(R_q)^k$ as well.

If b is uniform, the first sample yields $q^{-1}s_0 \bmod R^\vee$ uniform over $q^{-1}R^\vee/R^\vee$. For the other samples, $b + q^{-1}(s_0 \cdot a'') \bmod R^\vee$ is uniform over \mathbb{T}_{R^\vee} and independent of $\mathbf{U}(a''|\mathbf{a})$ but also independent from the first sample because b masks $q^{-1}(s_0 \cdot a'')$.

If $b = q^{-1}\langle \mathbf{a}, \mathbf{s} \rangle + e \bmod R^\vee$ for some uniform $\mathbf{s} \in (R_q^\vee)^{k-1}$ and $e \leftarrow \psi$ for some $\psi \leftarrow \Upsilon$, then $q^{-1}s_0 = q^{-1}\langle \mathbf{e}_1, (s_0|\mathbf{s}) \rangle = q^{-1}\langle \mathbf{U}\mathbf{e}_1, \mathbf{U}^{-T}(s_0|\mathbf{s}) \rangle = q^{-1}\langle \mathbf{a}', \mathbf{U}^{-T}(s_0|\mathbf{s}) \rangle$, where $\mathbf{e}_1 = [1, 0, \dots, 0]^T$. For the other samples, we have

$$\begin{aligned} b + q^{-1}(s_0 \cdot a'') \bmod R^\vee &= q^{-1}\langle \mathbf{a}, \mathbf{s} \rangle + q^{-1}(s_0 \cdot a'') + e \bmod R^\vee \\ &= q^{-1}\langle (a''|\mathbf{a}), (s_0|\mathbf{s}) \rangle + e \bmod R^\vee \\ &= q^{-1}\langle \mathbf{U}(a''|\mathbf{a}), \mathbf{U}^{-T}(s_0|\mathbf{s}) \rangle + e \bmod R^\vee. \end{aligned}$$

Note that $(s_0|\mathbf{s})$ is uniform over $(R_q^\vee)^k$ so $\mathbf{U}^{-T}(s_0|\mathbf{s})$ is also uniform over $(R_q^\vee)^k$ because \mathbf{U}^{-T} is invertible in R_q . Therefore the reduction outputs samples according to first-is-errorless M-LWE with secret $\mathbf{s}' = \mathbf{U}^{-T}(s_0|\mathbf{s})$. ■

2.3.2 Extended M-LWE

We now define the module version of the extended variant of LWE introduced in [BLP⁺13], where the adversary is allowed a hint on the errors. As opposed to [AA16], who also define such a module version, we allow for multiple secrets and one single hint vector \mathbf{z} , as required by our final reduction of Lemma 2.5.

Definition 2.2 (Extended M-LWE): Let $m, q, k, t, n \in \mathbb{N}$ and K be a number field of degree n with R its ring of integers. Let $\mathcal{Z} \subseteq (R^\vee)^m$ and ψ be a discrete distribution over $q^{-1}(R^\vee)^m$. The ext-M-LWE $_{n,k,q,\psi,\mathcal{Z}}^{m,t}$ problem is as follows: A PPT adversary first chooses $\mathbf{z} \in \mathcal{Z}$ and then receives a tuple

$$(\mathbf{A}, (\mathbf{b}_j)_{j \in [t]}, (\langle \mathbf{e}_j, \mathbf{z} \rangle)_{j \in [t]}) \in (R_q)^{k \times m} \times ((q^{-1}R^\vee/R^\vee)^m)^t \times (q^{-1}R^\vee)^t.$$

Their goal is to distinguish between the following cases. On one side, \mathbf{A} is sampled uniformly over $(R_q)^{k \times m}$, and for all $j \in [t]$, $\mathbf{e}_j \in q^{-1}(R^\vee)^m$ are independent and identically distributed from ψ , and define $\mathbf{b}_j = q^{-1}\mathbf{A}^T \mathbf{s}_j + \mathbf{e}_j \bmod R^\vee$ for some uniformly cho-

sen $\mathbf{s}_j \in (R_q^\vee)^k$. On the other side, everything is identical except that the \mathbf{b}_j are sampled uniformly over $(q^{-1}R^\vee/R^\vee)^m$, independently from \mathbf{A} and the error vectors.

The parameter n defines the ring degree, q the modulus, k the module rank (i.e., the number of columns of the matrix \mathbf{A}), m the number of samples (i.e., the number of rows of the matrix \mathbf{A}) and t defines the number of given hints on independent noise vectors.

The set \mathcal{Z} represents the set of hints that can be given on the noise vector. The t hints are given in form of the inner-product of such a fixed *hint vector* $\mathbf{z} \in \mathcal{Z}$ and the corresponding noise vector \mathbf{e}_j for $j \in [t]$. Later, we are interested in the special set of vectors with coefficients in $\{0, \dots, \eta - 1\}$ for some $\eta \in \mathbb{N}$, thus in the case $\mathcal{Z} = (R_\eta^\vee)^m$. Especially the case of binary vectors (i.e., $\eta = 2$) is of interest in the improved reduction for bin-M-LWE.

For simplicity in what follows, for a matrix $\mathbf{A} \in R^{m \times m}$, we denote by $\mathbf{A}^\perp \in R^{m \times (m-1)}$ the submatrix of \mathbf{A} obtained by removing the leftmost column. Our reduction from the problem first-is-errorless M-LWE to ext-M-LWE in Lemma 2.3 requires the construction of a matrix $\mathbf{U}_z \in R^{m \times m}$, for all hint vectors $\mathbf{z} \in \mathcal{Z} = (R_\eta^\vee)^m$, satisfying several properties. This matrix allows us to transform samples from a first-is-errorless M-LWE challenger into samples that we can give to an oracle for ext-M-LWE. The largest singular value of its submatrix \mathbf{U}_z^\perp (when embedded with θ as defined in Section 1.3.2), controls the increase in the Gaussian parameter. We propose a construction for which we bound the largest singular value above by a quantity independent on \mathbf{z} , as needed in the reduction.

Lemma 2.2 (Construction of Matrix \mathbf{U}_z for Hint Vector \mathbf{z})

Let $\nu \in \mathbb{N}$ with prime-power factorization $\nu = \prod_j p_j^{e_j}$. Further, let K be the ν -th cyclotomic field of degree $n = \varphi(\nu)$, and R its ring of integers. Let $\eta \in \mathbb{N}$ and let $\mu = \prod_j p_j$ and q be a prime number such that $q \equiv 1 \pmod{\mu}$, $\text{ord}_\nu(q) = \nu/\mu$ and $q > ((\eta-1)\mathfrak{s}_1(\mu))^{\varphi(\mu)}$, where $\mathfrak{s}_1(\mu)$ denotes the largest singular value of the Vandermonde matrix of the μ -th cyclotomic field. Finally, let m be a positive integer, and $\mathcal{Z} = (R_\eta^\vee)^m$, and we recall the ring parameter $B = \max_{x \in R_\eta} \|\sigma(x)\|_\infty$. For all $\mathbf{z} \in \mathcal{Z}$, there is an efficiently computable matrix $\mathbf{U}_z \in R^{m \times m}$ that is invertible modulo qR and that verifies the following: \mathbf{z} is orthogonal to the columns of \mathbf{U}_z^\perp , and the largest singular value of $\theta(\mathbf{U}_z^\perp) \in \mathbb{C}^{m \times (m-1)n}$ is at most $2B$.

Proof: Recall from Section 1.1.5 that for any cyclotomic number field, there exists an element λ such that $\lambda R_\eta^\vee = R_\eta$. Let $\mathbf{z} \in \mathcal{Z}$ and denote $(\tilde{z}_j)_{j \in [m]} = \tilde{\mathbf{z}} = \lambda \mathbf{z} \in R_\eta^m$. First, we construct \mathbf{U}_z in the case where all the \tilde{z}_j are non-zero. To do so, we define $\mathbf{U}_z = \mathbf{A} + \mathbf{B}$ via the intermediate matrices \mathbf{A} and \mathbf{B} of $R^{m \times m}$, all unspecified entries being zeros:

$$\mathbf{U}_z = \begin{bmatrix} 1 & & & & & \\ & \tilde{z}_2 & & & & \\ & & \tilde{z}_1 & & & \\ & & & & & \\ & & & & & \\ & & & & & \tilde{z}_m \\ & & & & & & \tilde{z}_{m-1} \end{bmatrix} = \begin{bmatrix} 1 & & & & & \\ & \tilde{z}_1 & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \\ & & & & & \tilde{z}_{m-1} \end{bmatrix} + \begin{bmatrix} 0 & & & & & \\ & \tilde{z}_2 & & & & \\ & & \tilde{z}_1 & & & \\ & & & & & \\ & & & & & \\ & & & & & \tilde{z}_m \\ & & & & & & \tilde{z}_{m-1} \\ & & & & & & & 0 \end{bmatrix}$$

$\mathbf{U}_z^\perp \qquad \qquad \mathbf{A}^\perp \qquad \qquad \mathbf{B}^\perp$

The matrix $\mathbf{U}_{\mathbf{z}}$ is invertible modulo qR only if all the \tilde{z}_j (except \tilde{z}_m) are in R_q^\times . Yet, since they are all non-zero elements of R_η , we have that for all $j \in [m]$, $\|\tau(\tilde{z}_j)\|_\infty \leq \eta - 1$, where τ is the coefficient embedding. By Lemma 1.2, since q verifies the algebraic conditions taking all $f_j = 1$ and $q^{1/\varphi(\mu)}/\mathfrak{s}_1(\mu) > (\eta - 1)$, all the \tilde{z}_j are in R_q^\times .

By construction, the last $m-1$ columns of $\mathbf{U}_{\mathbf{z}}$ are orthogonal to $\tilde{\mathbf{z}}$. Let $\mathbf{U}_{\mathbf{z}}^\perp$ be the submatrix of $\mathbf{U}_{\mathbf{z}}$ obtained by removing the leftmost column as shown above. Since θ is a ring homomorphism, we have $\theta(\mathbf{U}_{\mathbf{z}}^\perp) = \theta(\mathbf{A}^\perp) + \theta(\mathbf{B}^\perp)$. We now need to bound the spectral norms (which equal their largest singular values) of these two matrices, and use the triangle inequality to conclude. For any vector $\mathbf{x} \in \mathbb{C}^{(m-1)n}$, we have that

$$\left\| \theta(\mathbf{A}^\perp) \mathbf{x} \right\|_2 = \sqrt{\sum_{j \in [m-1]} \sum_{k \in [n]} |\sigma_k(\tilde{z}_j)|^2 |x_{k+n(j-1)}|^2} \leq B \|\mathbf{x}\|_2,$$

because each \tilde{z}_k is in R_η . This yields $\left\| \theta(\mathbf{A}^\perp) \right\|_2 \leq B$. A similar calculation on \mathbf{B}^\perp leads to $\left\| \theta(\mathbf{B}^\perp) \right\|_2 \leq B$, thus resulting in $\left\| \theta(\mathbf{U}_{\mathbf{z}}^\perp) \right\|_2 \leq 2B$.

Now assume that $\tilde{z}_{j_0}, \dots, \tilde{z}_m$ are zeros for some j_0 in $[m]$. If the zeros do not appear last in the vector $\tilde{\mathbf{z}}$, we can replace $\tilde{\mathbf{z}}$ with $\mathbf{S}\tilde{\mathbf{z}}$, where $\mathbf{S} \in R^{m \times m}$ swaps the coordinates of $\tilde{\mathbf{z}}$ so that the zeros appear last. Since \mathbf{S} is unitary, it preserves the singular values as well as invertibility. Then, the construction remains the same except that the $\tilde{z}_{j_0}, \dots, \tilde{z}_m$ on the diagonal are replaced by 1. The orthogonality is preserved, and $\left\| \theta(\mathbf{U}_{\mathbf{z}}^\perp) \right\|_2$ can still be bounded above by $2B$. ■

Notice that when the ring is of degree 1 and $\eta = 2$ (binary hint vectors), the constructions in the different cases match the ones from [BLP⁺13, Claim 4.6]. So do the singular values as $B \leq (\eta - 1)n = 1$ by Lemma 1.1. Also, the construction differs from the notion of quality in [AA16] due to the discrepancies between the two definitions of ext-M-LWE. The following lemma shows that the extended variant of M-LWE with one hint ($t = 1$) is at least as hard as the first-is-errorless variant of M-LWE, for appropriate parameters.

Lemma 2.3 (first-is-errorless M-LWE to ext-M-LWE, adapted from [BLP⁺13, Lem. 4.7])

Given $\eta \in \mathbb{N}$ and let $\nu \in \mathbb{N}$ with prime-power-factorization $\nu = \prod_j p_j^{e_j}$. Further, let K be the ν -th cyclotomic field of degree $n = \varphi(\nu)$, and R its ring of integers. Let $\mu = \prod_j p_j$ and q be a prime such that $q \equiv 1 \pmod{\mu}$, $\text{ord}_\nu(q) = \nu/\mu$ and $q > ((\eta - 1)\mathfrak{s}_1(\mu))^{\varphi(\mu)}$, where $\mathfrak{s}_1(\mu)$ denotes the largest singular value of the Vandermonde matrix of the μ -th cyclotomic field. Let m, k positive integers, $\mathcal{Z} = (R_\eta^\vee)^m$, $\varepsilon \in (0, 1/2)$ and $\alpha \geq q^{-1} \sqrt{\ln(2mn(1 + 1/\varepsilon))}/\pi$. Then, there is a PPT reduction from first-is-errorless M-LWE $_{n,k,q,\psi}^m$ to ext-M-LWE $_{n,k,q,\psi',\mathcal{Z}}^m$, where $\psi = D_\alpha$ and $\psi' = \mathcal{D}_{q^{-1}(R^\vee)^m, \alpha \sqrt{4B^2+1}}$, where $B = \max_{x \in R_\eta} \|\sigma(x)\|_\infty$. The reduction reduces the advantage by at most $33\varepsilon/2$.

Proof: We begin with showing that the Gaussian width α is at least as large as the smoothing parameter of $q^{-1}(R^\vee)^m$. Using Lemma 1.4, we have $\lambda_1^\infty(R) \geq N(R)^{1/n} = 1$. So, using the fact

that $(q\Lambda)^* = q^{-1}\Lambda^*$, we have

$$\lambda_1^\infty((q^{-1}(R^\vee)^m)^*) = \lambda_1^\infty(q((R^\vee)^m)^*) = q\lambda_1^\infty(((R^\vee)^m)^*) = q\lambda_1^\infty(R) \geq q,$$

which together with Lemma 1.7 yields $\alpha \geq q^{-1}\sqrt{\ln(2mn(1+1/\varepsilon))/\pi} \geq \eta_\varepsilon(q^{-1}(R^\vee)^m)$.

Now, we show the reduction. Assume we have access to an oracle \mathcal{O} for ext-M-LWE $_{n,k,q,\psi',\mathcal{Z}}^m$, where $\psi' = \mathcal{D}_{q^{-1}(R^\vee)^m, \alpha\sqrt{4B^2+1}}$. We request m samples from the first-is-errorless challenger, resulting in

$$(\mathbf{A}, \mathbf{b}) \in (R_q)^{k \times m} \times ((q^{-1}R^\vee/R^\vee) \times \mathbb{T}_{R^\vee}^{m-1}).$$

Assume we need to provide samples to \mathcal{O} for some $\mathbf{z} \in \mathcal{Z}$. By Lemma 2.2 we can efficiently compute a matrix $\mathbf{U}_\mathbf{z} \in R^{m \times m}$ that is invertible modulo qR , such that its submatrix $\mathbf{U}_\mathbf{z}^\perp$ is orthogonal to \mathbf{z} , and that $\theta(\mathbf{U}_\mathbf{z}^\perp)$ has largest singular value less than $2B$. The reduction first samples $\mathbf{f} \in K_{\mathbb{R}}^m$ from the continuous Gaussian distribution of covariance matrix $\alpha^2(4B^2\mathbf{I}_{mn} - \mathbb{H}^\dagger\theta(\mathbf{U}_\mathbf{z}^\perp)\theta(\mathbf{U}_\mathbf{z}^\perp)^\dagger\mathbb{H}) \in \mathbb{R}^{mn \times mn}$, where \mathbb{H} is defined as in Section 1.3.2. Note that \mathbb{H} is unitary and therefore preserves the largest singular value. The reduction then computes $\mathbf{b}' = \mathbf{U}_\mathbf{z}\mathbf{b} + \mathbf{f}$ and samples \mathbf{c} from $\mathcal{D}_{q^{-1}(R^\vee)^m - \mathbf{b}', \alpha}$, and finally gives the following to \mathcal{O}

$$(\mathbf{A}' = \mathbf{A}\mathbf{U}_\mathbf{z}^T, \mathbf{b}' + \mathbf{c} \bmod R^\vee, \langle \mathbf{z}, \mathbf{f} + \mathbf{c} \rangle).$$

Note that this tuple is in $(R_q)^{k \times m} \times (q^{-1}R^\vee/R^\vee)^m \times q^{-1}R^\vee$, as required.

We now prove correctness. First, consider the case where \mathbf{A} is uniformly random over $R_q^{k \times m}$ and $\mathbf{b} = q^{-1}\mathbf{A}^T\mathbf{s} + \mathbf{e} \bmod R^\vee$ for some uniform $\mathbf{s} \in (R_q^\vee)^k$, and \mathbf{e} sampled from $\{0\} \times D_\alpha^{m-1}$ where $\{0\}$ denotes the distribution that is deterministically 0. Since $\mathbf{U}_\mathbf{z}$ is invertible modulo qR , $\mathbf{A}' = \mathbf{A}\mathbf{U}_\mathbf{z}^T$ is also uniform over $(R_q)^{k \times m}$ as required. From now on we condition on an arbitrary \mathbf{A}' and analyze the distribution of the remaining components. We have

$$\begin{aligned} \mathbf{b}' &= q^{-1}\mathbf{U}_\mathbf{z}\mathbf{A}^T\mathbf{s} + \mathbf{U}_\mathbf{z}\mathbf{e} + \mathbf{f} \\ &= q^{-1}(\mathbf{A}')^T\mathbf{s} + \mathbf{U}_\mathbf{z}\mathbf{e} + \mathbf{f}. \end{aligned}$$

Since the first coefficient of \mathbf{e} is deterministically 0 the first column is ignored in the covariance matrix, and then $\mathbf{U}_\mathbf{z}\mathbf{e}$ is distributed as the Gaussian over $K_{\mathbb{R}}^m$ of covariance matrix $\alpha^2\mathbb{H}^\dagger\theta(\mathbf{U}_\mathbf{z}^\perp)\theta(\mathbf{U}_\mathbf{z}^\perp)^\dagger\mathbb{H}$ by Lemma 1.13. Hence the vector $\mathbf{U}_\mathbf{z}\mathbf{e} + \mathbf{f}$ is distributed as the Gaussian over $K_{\mathbb{R}}^m$ of covariance matrix $\alpha^2\mathbb{H}^\dagger\theta(\mathbf{U}_\mathbf{z}^\perp)\theta(\mathbf{U}_\mathbf{z}^\perp)^\dagger\mathbb{H} + \alpha^2((2B)^2\mathbf{I}_{mn} - \mathbb{H}^\dagger\theta(\mathbf{U}_\mathbf{z}^\perp)\theta(\mathbf{U}_\mathbf{z}^\perp)^\dagger\mathbb{H})$ which is identical to $D_{\alpha 2B}^m$. Since $q^{-1}(\mathbf{A}')^T\mathbf{s} \in q^{-1}(R^\vee)^m$, the coset $q^{-1}(R^\vee)^m - \mathbf{b}'$ is the same as $q^{-1}(R^\vee)^m - (\mathbf{U}_\mathbf{z}\mathbf{e} + \mathbf{f})$, which yields that \mathbf{c} can be seen as being sampled from $\mathcal{D}_{q^{-1}(R^\vee)^m - (\mathbf{U}_\mathbf{z}\mathbf{e} + \mathbf{f}), \alpha}$. By the remark made at the beginning of the proof, we have $\alpha \geq \eta_\varepsilon(q^{-1}(R^\vee)^m)$, so by Lemma 1.11, the distribution of $\mathbf{U}_\mathbf{z}\mathbf{e} + \mathbf{f} + \mathbf{c}$ is within statistical distance 8ε of $\mathcal{D}_{q^{-1}(R^\vee)^m, \alpha\sqrt{4B^2+1}}$, which shows that the second component is correctly distributed up to 8ε . Note that $\mathbf{U}_\mathbf{z}\mathbf{e} = \sum_{i \in [m]} e_i \cdot \mathbf{u}_i$ is in the space spanned by the columns of $\mathbf{U}_\mathbf{z}^\perp$ because $e_1 = 0$. This yields $\langle \mathbf{z}, \mathbf{U}_\mathbf{z}\mathbf{e} \rangle = 0$ as \mathbf{z} is orthogonal to the columns of $\mathbf{U}_\mathbf{z}^\perp$. This proves that the third component equals $\langle \mathbf{z}, \mathbf{U}_\mathbf{z}\mathbf{e} + \mathbf{f} + \mathbf{c} \rangle$ and is thus correctly distributed.

Now consider the case where both \mathbf{A} and \mathbf{b} are uniform. First, observe that $\alpha \geq \eta_\varepsilon(q^{-1}(R^\vee)^m)$ and therefore by Lemma 1.9, the distribution of (\mathbf{A}, \mathbf{b}) is within statistical distance $\varepsilon/2$ of the distribution of $(\mathbf{A}, \mathbf{e}' + \mathbf{e})$ where $\mathbf{e}' \in (q^{-1}R^\vee/R^\vee)^m$ is uniform and \mathbf{e} is distributed from $\{0\} \times D_\alpha^{m-1}$. So we can assume our input is $(\mathbf{A}, \mathbf{e}' + \mathbf{e})$. \mathbf{A}' is uniform as before, and clearly independent of the other two components. Moreover, since $\mathbf{b}' = \mathbf{U}_\mathbf{z}\mathbf{e}' + \mathbf{U}_\mathbf{z}\mathbf{e} + \mathbf{f}$

and $\mathbf{U}_z \mathbf{e}' \in q^{-1}(R^\vee)^m$, then the coset $q^{-1}(R^\vee)^m - \mathbf{b}'$ is identical to $q^{-1}(R^\vee)^m - (\mathbf{U}_z \mathbf{e} + \mathbf{f})$. For the same reasons as above, $\mathbf{U}_z \mathbf{e} + \mathbf{f} + \mathbf{c}$ is distributed as $\mathcal{D}_{q^{-1}(R^\vee)^m, \alpha\sqrt{4B^2+1}}$ within statistical distance of at most 8ε , and in particular independent of \mathbf{e}' . So the third component is correctly distributed because once again $\langle \mathbf{z}, \mathbf{U}_z \mathbf{e} \rangle = 0$. Finally, since \mathbf{e}' is independent of the first and third components, and that $\mathbf{U}_z \mathbf{e}'$ is uniform over $(q^{-1}R^\vee/R^\vee)^m$ as \mathbf{U}_z is invertible modulo qR , it yields that the second component is uniform and independent of the other ones as required. \blacksquare

🔗 Example 2.1 (Instantiation in Power-of-2 Cyclotomics)

The condition on the modulus q in Lemma 2.2 and Lemma 2.3 stems from the invertibility result from Lyubashevsky and Seiler [LS18], as stated in Lemma 1.2. It can be simplified in the power-of-2 case, see Example 1.4, where it is conditioned on the number $\kappa > 1$ of splitting factors of the defining polynomial $x^n + 1$ in $\mathbb{Z}_q[x]$. Choosing κ as a power of two less than $n = 2^k$, implies that q has to be a prime congruent to $2\kappa + 1$ modulo 4κ . The invertibility condition then becomes $0 < \|\tau(y)\|_\infty < q^{1/\kappa}/\sqrt{\kappa}$ for any y in R_q . The upper bound is decreasing with κ so the smaller κ , the more invertible elements. The smallest choice for κ is $\kappa = 2$, which leads to choosing a prime $q = 5 \pmod{8}$. In our context, having $q^{1/2}/\sqrt{2} > (\eta - 1)$ is sufficient as our elements have coefficients at most $\eta - 1$. For the special case of $\eta = 2$, this leads to $q > 2$ which is subsumed by $q = 5 \pmod{8}$.

Finally, we use a standard hybrid argument to show that ext-M-LWE with t hints is at least as hard as ext-M-LWE with 1 hint, at the expense of reducing the advantage by a multiplicative factor of t .

Lemma 2.4 (Adapted from [BLP⁺13, Lem. 4.8])

Let $n, k, m, q, t \in \mathbb{N}$ such that $t \leq \text{poly}(n)$ and K be a number field of degree n with R its ring of integers. Let ψ be a discrete distribution over $q^{-1}(R^\vee)^m$, and $\mathcal{Z} \subseteq (R^\vee)^m$. There is a reduction from ext-M-LWE $_{n,k,q,\psi,\mathcal{Z}}^m$ to ext-M-LWE $_{n,k,q,\psi,\mathcal{Z}}^{m,t}$ that reduces the advantage by a multiplicative factor of t .

Proof: Let \mathcal{O} be an oracle for ext-M-LWE $_{n,k,q,\psi,\mathcal{Z}}^{m,t}$. For each $k \in \{0, \dots, t\}$, we denote by \mathcal{H}_i the hybrid distribution defined as

$$(\mathbf{A}, (\mathbf{b}_1, \dots, \mathbf{b}_k, \mathbf{u}_{k+1}, \dots, \mathbf{u}_t), (\langle \mathbf{e}_j, \mathbf{z} \rangle)_{j \in [t]}),$$

where $\mathbf{A} \leftarrow U(R_q^{k \times m})$, the $\mathbf{u}_j \leftarrow U((q^{-1}R^\vee/R^\vee)^m)$, the $\mathbf{e}_j \leftarrow \psi$, and $\mathbf{b}_j = q^{-1}\mathbf{A}^T \mathbf{s}_j + \mathbf{e}_j \pmod{R^\vee}$ for $\mathbf{s}_j \leftarrow U((R_q^\vee)^k)$ for every $j \in [t]$. By definition, we have $\text{Adv}[\mathcal{O}] = |\Pr(\mathcal{O}(\mathcal{H}_t)) - \Pr(\mathcal{O}(\mathcal{H}_0))|$. The reduction \mathcal{A} works as follows.

1. Sample \mathbf{z} uniformly from \mathcal{Z} and get $(\mathbf{A}, \mathbf{b}, y = \langle \mathbf{e}, \mathbf{z} \rangle)$ as input of ext-M-LWE $_{n,k,q,\psi,\mathcal{Z}}^m$.
2. Sample k^* uniformly from $[t]$. This defines where to embed the challenge.
3. Sample $\mathbf{s}_1, \dots, \mathbf{s}_{k^*-1}$ uniformly from $(R_q^\vee)^k$, $\mathbf{e}_1, \dots, \mathbf{e}_{k^*-1}, \mathbf{e}_{k^*+1}, \dots, \mathbf{e}_t$ from ψ and fi-

- nally $\mathbf{u}_{k^*+1}, \dots, \mathbf{u}_t$ uniformly from $(q^{-1}R^\vee/R^\vee)^m$.
4. Compute $\mathbf{b}_j = q^{-1}\mathbf{A}^T \mathbf{s}_j + \mathbf{e}_j \bmod R^\vee$ for all $j \in [k^* - 1]$, and $y_j = \langle \mathbf{e}_j, \mathbf{z} \rangle$ for all $j \in [t] \setminus \{k^*\}$.
 5. Define $(\mathbf{b}'_j)_{j \in [t]}$ as $(\mathbf{b}_1, \dots, \mathbf{b}_{k^*-1}, \mathbf{b}, \mathbf{u}_{k^*+1}, \dots, \mathbf{u}_t)$. Then call the oracle \mathcal{O} on input $(\mathbf{A}, (\mathbf{b}'_j)_{j \in [t]}, (y_1, \dots, y_{k^*-1}, y, y_{k^*+1}, \dots, y_t))$, and return the same output as \mathcal{O} .

If \mathbf{b} is uniform, then the distribution in 5. is exactly \mathcal{H}_{k^*-1} whereas if \mathbf{b} is M-LWE, then the distribution is \mathcal{H}_{k^*} . By a standard hybrid argument, the oracle can distinguish between the two for some k^* if it can distinguish between \mathcal{H}_0 and \mathcal{H}_t . So the output is correct over the randomness of k^* . Since k^* is uniformly chosen we have

$$\begin{aligned} \text{Adv}[\mathcal{A}] &= |\Pr(\mathcal{A}(\mathbf{b} \text{ M-LWE})) - \Pr(\mathcal{A}(\mathbf{b} \text{ uniform}))| \\ &= \left| \sum_{k^* \in [t]} \frac{1}{t} \Pr(\mathcal{A}(\mathcal{H}_{k^*})) - \sum_{k^* \in [t]} \frac{1}{t} \Pr(\mathcal{A}(\mathcal{H}_{k^*-1})) \right| \\ &= \frac{1}{t} \text{Adv}[\mathcal{O}] \end{aligned}$$

2.3.3 Reduction to Binary M-LWE

We now provide the final step of the overall reduction, by reducing to the binary secret version of M-LWE using a sequence of hybrids. The idea is to use the set \mathcal{Z} of the ext-M-LWE problem as our set of secrets.

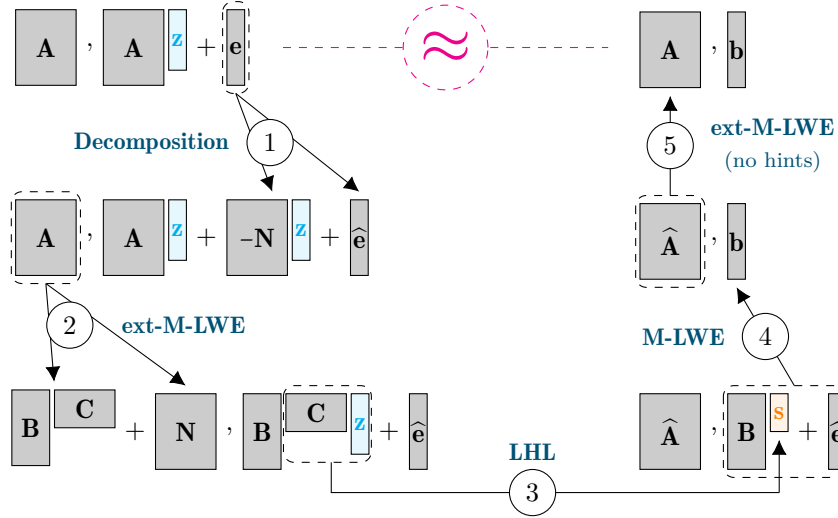


Figure 2.4: Overview of the proof of Lemma 2.5.

To facilitate understanding, we start by illustrating the high level idea of the proof of Lemma 2.5 in Figure 2.4 (where the Latex/tikz credits go to my co-author Corentin Jeudy). Given an instance of bin-M-LWE by $(\mathbf{A}, \mathbf{A} \cdot \mathbf{z} + \mathbf{e})$, our goal is to show that it is computationally

indistinguishable from (\mathbf{A}, \mathbf{b}) , where \mathbf{b} is a uniform random vector. To do so, we first decompose the error vector \mathbf{e} into $-\mathbf{N} \cdot \mathbf{z} + \widehat{\mathbf{e}}$, by using properties of Gaussian distributions. We then make use of a similar lossy argument as for the simple reduction of Section 2.2 by replacing the random matrix \mathbf{A} by the lossy matrix $\widehat{\mathbf{A}} = \mathbf{B} \cdot \mathbf{C} + \mathbf{N}$. As opposed to the proof from Section 2.2, we can't simply argue with the hardness of multiple-secrets M-LWE as the second part of the sample depends on the noise matrix \mathbf{N} . This is the motivation of the introduction of the extended variant of M-LWE, where we allow for additional information with respect to the noise. We then use the same Leftover Hash Lemma as before to replace the product $\mathbf{C} \cdot \mathbf{z}$ by a uniform random vector \mathbf{s} . Assuming the hardness of M-LWE, the product $\mathbf{B} \cdot \mathbf{s} + \widehat{\mathbf{e}}$ is computationally indistinguishable from a uniform vector \mathbf{b} . We conclude the proof by re-replacing the lossy matrix $\widehat{\mathbf{A}}$ by the original uniform matrix \mathbf{A} .

Lemma 2.5 (ext-M-LWE to bin-M-LWE, adapted from [BLP⁺13, Lem. 4.9])

Let $K = \mathbb{Q}(\zeta)$ be a number field of degree n with $R = \mathbb{Z}[\zeta]$ its ring of integers. Let q be a prime modulus and let k, m, d be positive integers such that $d \geq k \log_2 q + \omega(\log_2 n)$. Further, let $\varepsilon \in (0, 1/2)$ and $\alpha, \gamma, \beta, \delta$ be positive reals such that $\alpha \geq q^{-1} \sqrt{2 \ln(2nd(1+1/\varepsilon))} / \pi$, $\gamma = \alpha B \sqrt{d}$, $\beta = \alpha B \sqrt{2d}$, where $B = \max_{x \in R_2} \|\sigma(x)\|_\infty$, and $\delta = \frac{1}{2} \sqrt{(1 + q^k/2^d)^n - 1}$. Then there is a reduction from ext-M-LWE $_{n,k,q,\psi,(R_2^\vee)^d}^{d,m}$, M-LWE $_{n,k,q,D_\gamma}^m$ and ext-M-LWE $_{n,k,q,\psi,\{0\}^d}^{d,m}$ with $\psi = \mathcal{D}_{q^{-1}(R^\vee)^m, \alpha}$ to bin-M-LWE $_{n,d,q,\Psi \leq \beta}^m$, such that if $\mathcal{B}_1, \mathcal{B}_2$ and \mathcal{B}_3 are the algorithms obtained by applying these hybrids to an algorithm \mathcal{A} , then

$$\text{Adv}[\mathcal{A}] \leq \text{Adv}[\mathcal{B}_1] + \text{Adv}[\mathcal{B}_2] + \text{Adv}[\mathcal{B}_3] + 2m\varepsilon + \delta.$$

The problem ext-M-LWE $_{n,k,q,\psi,\{0\}^d}^{d,m}$ with $\psi = \mathcal{D}_{q^{-1}(R^\vee)^m, \alpha}$ mentioned in the lemma statement is trivially harder than ext-M-LWE $_{n,k,q,\psi,(R_2^\vee)^d}^{d,m}$, that is also why it is not specified in Figure 2.3.

Proof: For $x \in R^\vee$, we denote $\tilde{x} = \lambda x \in R$ as before, where λ induces a mapping from R^\vee to R as described in Section 1.1.5. We extend this notation to vectors and matrices in the obvious way. Given a bin-M-LWE $_{n,d,q,\Psi \leq \beta}^m$ sample $(\mathbf{A}, \mathbf{b} = q^{-1} \mathbf{A}^T \mathbf{z} + \mathbf{e} \bmod R^\vee)$, with $\mathbf{A} \leftarrow U((R_q)^{d \times m})$, $\mathbf{z} \leftarrow U((R_2^\vee)^d)$ and $\mathbf{e} \in K_{\mathbb{R}}^m$ sampled from the continuous Gaussian $D_{\mathbf{r}}^m$ with parameter vector \mathbf{r} with $r_j^2 = \gamma^2 + \alpha^2 \sum_i |\sigma_j(\tilde{z}_i)|^2$. Yet, we have $\|\mathbf{r}\|_\infty = \sqrt{\gamma^2 + \alpha^2 \|\tilde{\mathbf{z}}\|_{2,\infty}^2}$, as well as $\|\tilde{\mathbf{z}}\|_{2,\infty}^2 \leq \sum_{i \in [d]} \|\sigma(\tilde{z}_i)\|_\infty^2$. Recalling for the parameter $B = \max_{x \in R_2} \|\sigma(x)\|_\infty$, that can be bounded above by n for cyclotomics by Lemma 1.1 (for $\eta = 2$), we get $\|\mathbf{r}\|_\infty \leq \sqrt{\gamma^2 + B^2 d \alpha^2} = B \sqrt{2d} \alpha = \beta$. Our goal is to show that (\mathbf{A}, \mathbf{b}) is computationally indistinguishable from uniform. To do so, we define different hybrid distributions:

- H_0 : (\mathbf{A}, \mathbf{b}) as in bin-M-LWE $_{n,d,q,\Psi \leq \beta}^m$,
- H_1 : $(\mathbf{A}, q^{-1} \mathbf{A}^T \mathbf{z} - \lambda \mathbf{N}^T \mathbf{z} + \widehat{\mathbf{e}} \bmod R^\vee)$, where $\mathbf{N} \leftarrow \mathcal{D}_{q^{-1}R^\vee, \alpha}^{d \times m}$ and $\widehat{\mathbf{e}} \leftarrow D_\gamma^m$,
- H_2 : $(\widehat{\mathbf{A}}, q^{-1} \widehat{\mathbf{A}}^T \mathbf{z} - \lambda \mathbf{N}^T \mathbf{z} + \widehat{\mathbf{e}} \bmod R^\vee) = (\widehat{\mathbf{A}}, q^{-1} (\lambda \mathbf{B})^T \mathbf{C} \mathbf{z} + \widehat{\mathbf{e}} \bmod R^\vee)$ where \mathbf{B} is uniformly sampled over $(R_q^\vee)^{k \times m}$, \mathbf{C} uniformly sampled over $R_q^{k \times d}$ and $\widehat{\mathbf{A}} = \lambda q (q^{-1} \mathbf{C}^T \mathbf{B} + \mathbf{N} \bmod R^\vee)$,

- H_3 : $(\widehat{\mathbf{A}}, q^{-1}\widetilde{\mathbf{B}}^T \mathbf{s} + \widehat{\mathbf{e}} \bmod R^\vee)$, with $\widehat{\mathbf{A}}, \widehat{\mathbf{e}}$ and \mathbf{B} as above and where $\widetilde{\mathbf{B}} = \lambda \mathbf{B} \in R_q^{k \times m}$, and \mathbf{s} is uniform over $(R_q^\vee)^k$,
- H_4 : $(\widehat{\mathbf{A}}, \mathbf{u})$, with $\widehat{\mathbf{A}}$ as above and $\mathbf{u} \leftarrow U(\mathbb{T}_{R^\vee}^m)$,
- H_5 : $(\mathbf{A}, \mathbf{b}) \leftarrow U((R_q)^{d \times m} \times \mathbb{T}_{R^\vee}^m)$.

From H_0 to H_1 . By looking at each component of the vectors we claim that $\Delta([\mathbf{-N}^T \widetilde{\mathbf{z}} + \widehat{\mathbf{e}}]_i, \mathbf{e}_i) \leq 2\varepsilon$. Indeed, $(1/\alpha^2 + \|\widetilde{\mathbf{z}}\|_{2,\infty}^2/\gamma^2)^{-1/2} \geq \alpha/\sqrt{2}$ and $\alpha/\sqrt{2} \geq \eta_\varepsilon(q^{-1}(R^\vee)^d)$ as explained for Lemma 2.3. If $\mathbf{n}_i \in q^{-1}(R^\vee)^d$ denotes the i -th column of \mathbf{N} , Lemma 1.16 yields the claim as $[\mathbf{-N}^T \widetilde{\mathbf{z}} + \widehat{\mathbf{e}}]_i = \langle \mathbf{n}_i, -\widetilde{\mathbf{z}} \rangle + \widehat{e}_i$, thus giving $\Delta([\mathbf{-N}^T \widetilde{\mathbf{z}} + \widehat{\mathbf{e}}]_i, \mathbf{e}_i) \leq 2m\varepsilon$. We can thus deduce

$$|\Pr(\mathcal{A}(H_0)) - \Pr(\mathcal{A}(H_1))| \leq 2m\varepsilon. \quad (2.1)$$

From H_1 to H_2 . We argue that a distinguisher between H_1 and H_2 can be used to derive an adversary \mathcal{B}_1 for ext-M-LWE $_{n,k,q,\psi,(R_2^\vee)^d}^{d,m}$, where $\psi = \mathcal{D}_{q^{-1}(R^\vee)^m, \alpha}$, with the same advantage. To do so, \mathcal{B}_1 transforms the samples from the challenger of the ext-M-LWE problem to samples defined in H_1 or the ones in H_2 depending on whether or not the received samples are uniform. In the uniform case, $(\mathbf{C}, (\lambda q)^{-1} \mathbf{A}, \mathbf{N}^T \mathbf{z})$ can be efficiently transformed into a sample from H_1 . Note that $(\lambda q)^{-1} \mathbf{A}$ indeed corresponds to the uniform case of ext-M-LWE, because \mathbf{A} is uniform over R_q and $(\lambda q)^{-1} R_q$ can be seen as $q^{-1} R^\vee / R^\vee$. In the other case, if we apply the same transformation to the ext-M-LWE sample $(\mathbf{C}, q^{-1} \mathbf{C}^T \mathbf{B} + \mathbf{N} \bmod R^\vee, \mathbf{N}^T \mathbf{z})$, it leads to a sample from H_2 . Hence, \mathcal{B}_1 is a distinguisher for ext-M-LWE $_{n,k,q,\psi,(R_2^\vee)^d}^{d,m}$, and

$$|\Pr(\mathcal{A}(H_1)) - \Pr(\mathcal{A}(H_2))| = \text{Adv}[\mathcal{B}_1]. \quad (2.2)$$

From H_2 to H_3 . By the Ring Leftover Hash Lemma stated in Lemma 1.23 (for $\eta = 2$), we have that $(\mathbf{C}, \mathbf{C}\widetilde{\mathbf{z}})$ is within statistical distance at most δ from $(\mathbf{C}, \widetilde{\mathbf{s}})$. Note that for cyclotomic fields, it yields $R = \mathbb{Z}[\zeta]$ and thus the index $[R : \mathbb{Z}[\zeta]] = 1$ is always coprime to q . By multiplying by λ^{-1} and using the fact that a function does not increase the statistical distance, we have that $\Delta((\mathbf{C}, \mathbf{C}\mathbf{z}), (\mathbf{C}, \mathbf{s})) \leq \delta$. Note that the condition $d \geq k \log_2 q + \omega(\log_2 n)$ implies $\delta \leq n^{-\omega(1)}$. This yields

$$|\Pr(\mathcal{A}(H_2)) - \Pr(\mathcal{A}(H_3))| \leq \delta. \quad (2.3)$$

From H_3 to H_4 . A distinguisher between H_3 and H_4 can be used to derive an adversary \mathcal{B}_2 for M-LWE $_{n,k,q,D_\gamma}^m$. For that, \mathcal{B}_2 applies the efficient transformation to the samples from the M-LWE challenger, which turns $(\widetilde{\mathbf{B}}, \mathbf{u})$ into a sample from H_4 in the uniform case, and $(\widetilde{\mathbf{B}}, q^{-1}\widetilde{\mathbf{B}}^T \mathbf{s} + \widehat{\mathbf{e}} \bmod R^\vee)$ into a sample from H_3 in the M-LWE case. Therefore, \mathcal{B}_2 is a distinguisher for M-LWE $_{n,k,m,q,D_\gamma}^m$ such that

$$|\Pr(\mathcal{A}(H_3)) - \Pr(\mathcal{A}(H_4))| = \text{Adv}[\mathcal{B}_2]. \quad (2.4)$$

From H_4 to H_5 . In the last hybrid, we change $\widehat{\mathbf{A}}$ back to uniform. With the same argument as for the second hybrid, we can construct an adversary \mathcal{B}_3 for ext-M-LWE $_{n,k,q,\psi,\{0\}^d}^{d,m}$ with $\psi = \mathcal{D}_{q^{-1}(R^\vee)^m, \alpha}$ (which corresponds to multiple-secret M-LWE without additional information on the error) based on a distinguisher between H_4 and H_5 . It transforms $(\mathbf{C}, (\lambda q)^{-1} \widehat{\mathbf{A}}, \mathbf{N}^T \mathbf{0})$ into

a sample from H_4 (M-LWE case) and $(\mathbf{C}, (\lambda q)^{-1}\mathbf{A}, \mathbf{N}^T\mathbf{0})$ into a sample from H_5 (uniform case). We then get

$$|\Pr(\mathcal{A}(H_4)) - \Pr(\mathcal{A}(H_5))| = \text{Adv}[\mathcal{B}_3]. \quad (2.5)$$

Putting Equations 2.1, 2.2, 2.3, 2.4 and 2.5 altogether yields the result. \blacksquare

2.4 Generalization to Larger Secrets

An interesting question to ask is if the results of this chapter generalize to other secret distributions, in particular distributions where the secret's coefficients are not necessarily chosen uniformly at random over $\{0, 1\}$ but over a slightly larger set $\{0, \dots, \eta - 1\}$ for some positive integer $\eta \ll q$. We denote the corresponding problem by η -M-LWE. Surprisingly, both reductions that we have seen in this chapter can be easily adapted to this case, as we explain in the following.

Generalization for Warm-up Reduction

For the warm-up reduction from Section 2.2, recall that the distribution of the binary secret \mathbf{s} intervenes twice in the proof of Theorem 2.1. First, when we apply the LHL over rings to replace the term $\mathbf{C}\mathbf{s}$ by some uniform secret \mathbf{s}' and second, when we use the noise flooding technique to replace the term $\mathbf{Z}\mathbf{s} + \mathbf{e}$ with a fresh error vector \mathbf{e}' .

Fortunately, the LHL (Lemma 1.23) is already general enough to apply in the case where the secret \mathbf{s} is sampled uniformly at random over $R_\eta = R/\eta R$. To ease readability, we keep the following discussion over the ring R , assuming that we already applied the scalar λ to go from R^\vee to R . As we are considering the search variants of M-LWE and η -M-LWE, we can apply the LHL with respect to the Rényi divergence. In this case, the Rényi divergence between $(\mathbf{C}, \mathbf{C}\mathbf{s})$ and $(\mathbf{C}, \mathbf{s}')$, where $\mathbf{C} \leftarrow U((R_q)^{k \times d})$, $\mathbf{s} \leftarrow U((R_\eta)^d)$ and $\mathbf{s}' \leftarrow U((R_q)^k)$, is bounded above by $(1 + q^k/\eta^d)^n$. Thus, in order to obtain a constant Rényi divergence between the hybrid distributions H_3 and H_4 in the proof of Theorem 2.1, we have to require

$$d \geq k \cdot \frac{\log_2 q}{\log_2 \eta} + O\left(\frac{\log_2 n}{\log_2 \eta}\right).$$

Concerning the noise flooding technique, used to bound the Rényi divergence between the hybrid distributions H_2 and H_3 in the proof of Theorem 2.1, it suffices to apply Lemma 1.14, which is already proven for general secret distributions. It implies a resulting error of $\beta \geq \alpha n^2 \cdot d(\eta - 1)\sqrt{m}$. This leads to the following generalization of Theorem 2.1.

Theorem 2.3 (Hardness of M-LWE With Small Secret)

Let K be a cyclotomic number field of degree n with R its ring of integers. Let $k, d, m, q, \eta \in \mathbb{N}$ such that q is prime and m polynomial in n . Further, let α and β be positive real numbers such that $\beta \geq (\eta - 1)\sqrt{m} \cdot n^2 d \cdot \alpha$. Let $\varepsilon \in \mathbb{R} \cap [0, 1/2)$ such that $\beta q \geq \eta_\varepsilon(R^\vee)$ and $\varepsilon = O(\frac{1}{m})$. Then, for any $d \geq k \cdot (\log_2 q / \log_2 \eta) + \Omega(\log_2 n / \log_2 \eta)$, there is a probabilistic polynomial-time reduction from $\text{M-SLWE}_{n,k,q,\mathcal{D}_{R^\vee,\beta q}}^m$ and $\text{M-LWE}_{n,k,q,\mathcal{D}_{R^\vee,\alpha q}}^{m,d}$ to η - $\text{M-SLWE}_{n,d,q,\mathcal{D}_{R^\vee,\beta q}}^m$.

To summarize, we can see a direct link between the required rank, secret distribution and error growth with respect to the parameter η . If η gets close to 2, the error growth (determined

by β) gets smaller, but the requirement on rank d must be slightly larger. On the other hand, when η grows, the error ratio grows with it, but allows for a slightly smaller rank d . Note that since d also appears in the noise growth, the increase factor of considering η instead of 2 is roughly $(\eta - 1)/\log_2 \eta$.

Generalization for Improved Reduction

For the improved reduction from Section 2.3, we recall that the binary secret \mathbf{z} intervenes twice in the proof of Lemma 2.5. As before, it appears when we apply the LHL over rings to replace the term $\mathbf{C}\mathbf{z}$ by some uniform secret \mathbf{s} . Additionally, the binary secret determines the hints that we give as an additional information in the extended version of M-LWE.

As explained above, the LHL (Lemma 1.23) is already general enough to apply for more general secrets \mathbf{z} , which are sampled uniformly at random over $R_\eta = R/\eta R$. In contrast to the discussion above, we are now considering the decision variants of M-LWE and η -M-LWE, and thus we have to apply the LHL with respect to the statistical distance. More concretely, the statistical distance between $(\mathbf{C}, \mathbf{C}\mathbf{z})$ and (\mathbf{C}, \mathbf{s}) , where $\mathbf{C} \leftarrow U((R_q)^{k \times d})$, $\mathbf{z} \leftarrow U((R_\eta)^d)$ and $\mathbf{s} \leftarrow U((R_q)^k)$, is bounded above by $\frac{1}{2} \sqrt{(1 + q^k/\eta^d)^n - 1}$. Thus, in order to obtain a negligibly small statistical distance between the hybrid distributions H_2 and H_3 in the proof of Lemma 2.5, we require

$$d \geq k \cdot \frac{\log_2 q}{\log_2 \eta} + \omega \left(\frac{\log_2 n}{\log_2 \eta} \right).$$

Concerning the impact of \mathbf{z} on the extended variant of M-LWE, we need to consider ext-M-LWE with a set of hints given by $\mathcal{Z} = (R_\eta^\vee)^m$ (instead of $\mathcal{Z} = (R_2^\vee)^m$). Fortunately, Lemma 2.2 and Lemma 2.3 are already proven for this set of hints. The impact of allowing for a larger hint is reflected in the bound B from Lemma 1.1, which is given by $(\eta-1)n$ for cyclotomic fields, and in the invertibility condition from Lemma 1.2, which now requires a q larger by some factor $(\eta-1)^{\varphi(\mu)}$. Overall, we obtain a resulting error of $\beta \geq \alpha \cdot n(\eta-1)\sqrt{2d}\sqrt{4n^2(\eta-1)^2+1}$. This leads to the following generalization of Theorem 2.2.

Theorem 2.4 (Improved Hardness of M-LWE With Small Secret)

Let $\nu \in \mathbb{N}$ with prime-power factorization $\nu = \prod_j p_j^{e_j}$. Further, let K be the ν -th cyclotomic field of degree $n = \varphi(\nu)$, and R its ring of integers. Set $\mu = \prod_j p_j$, let $\eta \in \mathbb{N}$ and let q be a prime number such that $q \equiv 1 \pmod{\mu}$, $\text{ord}_\nu(q) = \nu/\mu$ and $q > \max(2n, ((\eta-1)\mathfrak{s}_1(\mu))^{\varphi(\mu)})$, where $\mathfrak{s}_1(\mu)$ denotes the largest singular value of the Vandermonde matrix of the μ -th cyclotomic field. Further, let $k, d, m \in \mathbb{N}$ such that $d \geq (k+1)(\log_2 q/\log_2 \eta) + \omega(\log_2 n/\log_2 \eta)$, and $d \leq m \leq \text{poly}(n)$. Let $\alpha \geq q^{-1} \sqrt{\ln(2nd(1+1/\varepsilon))/\pi}$ and $\beta \geq \alpha \cdot n(\eta-1)\sqrt{2d}\sqrt{4n^2(\eta-1)^2+1}$. Then there is a reduction from $\text{M-LWE}_{n,k,q,D_\alpha}^m$ to $\eta\text{-M-LWE}_{n,d,q,\Psi \leq \beta}^m$, such that if \mathcal{A} solves the latter with advantage $\text{Adv}[\mathcal{A}]$, then there exists an algorithm \mathcal{B} that solves the former with advantage

$$\text{Adv}[\mathcal{B}] \geq \frac{1}{3m} \left(\text{Adv}[\mathcal{A}] - \frac{1}{2} \sqrt{\left(1 + \frac{q^{k+1}}{2^d}\right)^n - 1} \right) - \frac{37\varepsilon}{2}.$$

As before, we can see a direct link between the required rank, secret distribution and error growth with respect to the parameter η . For increasing η , we can allow for a slightly smaller rank d , but obtain a larger noise ratio. Note that since d also appears in the noise growth, the increase factor of considering η instead of 2 is roughly $(\eta - 1)^2/\sqrt{\log_2 \eta}$.

2.5 Choice of Embedding for Binary Secrets

As mentioned in the introduction, the variant of M-LWE using a binary secret requires the choice of an embedding in which the secret is binary. As praised in [LPR10, LPR13], the canonical embedding has nice algebraic and geometric properties that make it a good choice of embedding. However, in this section, we justify our choice of the coefficient embedding, by analyzing the set of secrets that are binary in the canonical embedding in the case of power-of-2 cyclotomics.

The conjugation symmetry of the canonical embedding first restricts the choice of secrets to $(\sigma^{-1}(\{0, 1\}^n \cap H))^d$, where n denotes the ring degree, d the module rank and H the range of σ , see Section 1.1.1. In addition, the tightest worst-case to average-case reductions for M-LWE require \mathbf{s} to be taken from the dual $(R_q^\vee)^d$. However, σ^{-1} maps H to $K_{\mathbb{R}}$ but not necessarily to R or to R^\vee . We thus have to further restrict the set of secrets to

$$\mathcal{Z} = (R_q^\vee \cap \sigma^{-1}(\{0, \lambda^{-1}\}^n \cap H))^d,$$

where λ is such that $R^\vee = \lambda^{-1}R$, see Section 1.1.5. In the case of power-of-2 cyclotomics, it yields $\lambda = n$ and therefore $\lambda\mathcal{Z} = (R_q \cap \sigma^{-1}(\{0, 1\}^n \cap H))^d$.

2.5.1 Lagrange Basis

As opposed to R_2 which corresponds to binary vectors in the coefficient embedding, the power basis is not adapted to describe the set $\lambda\mathcal{Z}$. We thus introduce the Lagrange basis. Let $n = t_1 + 2t_2$. For $j \in [n]$, we denote by $\alpha_j = \sigma_j(\zeta)$ the j -th root of the defining polynomial f . Recall that we assume that α_j is real for $j \in [t_1]$, and that we have $\alpha_{t_1+j} = \overline{\alpha_{t_1+t_2+j}} \in \mathbb{C}$ for $j \in [t_2]$. Applying σ_j to an element $r = \sum_{k=0}^{n-1} r_k \zeta^k \in K_{\mathbb{R}}$ comes down to evaluating the polynomial $p_r = \sum_{k=0}^{n-1} r_k x^k$ at α_j . We use this polynomial interpretation to define elements of $K_{\mathbb{R}}$ that form a basis of $\sigma^{-1}(\{0, 1\}^n \cap H)$. Lagrange interpolation defines polynomials that map a set of distinct elements to 0 and 1. Since the α_j are distinct as f is irreducible, we can apply a similar method and define

$$L_k = \prod_{j \in [n] \setminus \{k\}} \frac{x - \alpha_j}{\alpha_k - \alpha_j},$$

for $k \in [t_1]$, which is real due to the conjugation symmetry of the roots. For $k \in \{t_1+1, \dots, t_1+t_2\}$, we define

$$L_k = \prod_{j \in [n] \setminus \{k\}} \frac{x - \alpha_j}{\alpha_k - \alpha_j} + \prod_{j \in [n] \setminus \{k+t_2\}} \frac{x - \alpha_j}{\alpha_{k+t_2} - \alpha_j} = 2\Re \left(\prod_{j \in [n] \setminus \{k\}} \frac{x - \alpha_j}{\alpha_k - \alpha_j} \right),$$

where $\Re(z)$ denotes the real component of a complex number z .

Hence the polynomials lie in $\mathbb{R}[x]$ and we have $L_k(\alpha_j) = \delta_{k,j}$ for $(k, j) \in [t_1] \times [n]$, and $L_k(\alpha_j) = \delta_{k,j} + \delta_{k+t_2,j}$ for $(k, j) \in \{t_1+1, \dots, t_1+t_2\} \times [n]$. Therefore, by defining the Lagrange basis \vec{l} with the corresponding $l_k \cong L_k(\zeta) \in K_{\mathbb{R}}$, we have linear independence and $\sigma^{-1}(\{0, 1\}^n \cap H) = \sum_{k \in [t_1+t_2]} \{0, 1\} \cdot l_k$, because $\sigma(l_k) = \mathbf{e}_k$ if $k \in [t_1]$ and $\sigma(l_k) = \mathbf{e}_k + \mathbf{e}_{k+t_2}$ if $k \in \{t_1+1, \dots, t_1+t_2\}$. As far as we are aware, this is the first time that the Lagrange basis is used in the setting of structured lattice-based cryptography. We now need to determine which of these combinations lie in R_q in order to properly define the set of secrets.

2.5.2 Lagrange Basis for Power-of-2 Cyclotomics

We now look at the Lagrange basis in the specific case where n is a power of 2.

Lemma 2.6 (Set of Binary Secrets for Power-of-2 Cyclotomics)

Let R be the cyclotomic ring of integers of degree $n = 2^\ell$ for $\ell \in \mathbb{N}$. Then, for any $q \in \mathbb{N}$, the set $R_q \cap \sigma^{-1}(\{0, 1\}^n \cap H)$ contains only 0 and 1.

Proof: Recall that in cyclotomic fields, we have $t_1 = 0$ and $t_2 = n/2$, see Section 1.1.2. We know that for power-of-2 cyclotomics the defining polynomial is given by $x^n + 1$ and therefore we can re-index the roots as $\alpha_j = \exp(i(2j+1)\pi/n)$, j now ranging from 0 to $n-1$. We can therefore study the complex product. We look at the constant coefficient of L_k , i.e., $A_k = L_k(0) = 2\Re\left(\prod_{0 \leq j < n, j \neq k} \frac{-\alpha_j}{\alpha_k - \alpha_j}\right)$. To ease notation, we write $j \neq k$ instead of $j \in \{0, \dots, n-1\} \setminus \{k\}$ for the product indexes. We first look at the product for a fixed $k \in \{0, \dots, n/2 - 1\}$.

$$\begin{aligned} \prod_{j \neq k} (\alpha_k - \alpha_j) &= \alpha_k^{n-1} \prod_{j \neq k} (1 - \alpha_j/\alpha_k) = -\alpha_k^{-1} \prod_{j \neq k} (1 - e^{i2\pi(j-k)/n}) \\ &= -\alpha_k^{-1} \prod_{l=1}^{n-1} (1 - e^{i2\pi l/n}), \end{aligned}$$

using the fact that $\alpha_k^n + 1 = 0$ and the circularity of the complex exponential. Yet, we also have $\prod_{l=0}^{n-1} (x - e^{i2\pi l/n}) = x^n - 1 = (x-1) \sum_{l=0}^{n-1} x^l$. By simplifying both sides by $x-1$ and then evaluating at 1, we have $\prod_{l=1}^{n-1} (1 - e^{i2\pi l/n}) = \sum_{l=0}^{n-1} 1^l = n$. The product of the numerators in the definition of A_k is $(-1)^{n-1} \overline{\alpha_k}$ because we can pair all of the roots α_j with their conjugates, which gives $\alpha_j \overline{\alpha_j} = |\alpha_j|^2 = 1$, except for $\overline{\alpha_k}$. Hence, $A_k = 2\Re(-\overline{\alpha_k}/(-n/\alpha_k))$ because n is even, which yields $A_k = \frac{2}{n}$. Now we take a subset $S \subseteq \{0, \dots, n/2 - 1\}$ and we study $\sum_{k \in S} L_k$. Note that the case of $S = \{0, \dots, n/2 - 1\}$ corresponds to adding all the Lagrange basis elements which results in 1, and the case $S = \emptyset$ results in 0 by convention. So we now assume that $0 < |S| < n/2$. The constant coefficient of $\sum_{k \in S} L_k$ is $2|S|/n \in (0, 1)$ and is therefore not an integer. Hence, $\sum_{k \in S} L_k \notin \mathbb{Z}[x]$ which means that the element $\sum_{k \in S} l_k$ is not in R nor R_q for any $q \geq 1$. It proves that the only binary combination of the Lagrange basis that are in R are 0 and 1, and the same conclusion is valid for R_q for any $q \geq 1$. \blacksquare

Hence to preserve the complexity of a brute force attack when comparing the two embeddings, the module rank would have to be increased by a factor n in the case where we take the canonical embedding to represent binary secrets. In this case, the (dual of the) secrets are from $\{0, 1\}^d$ and therefore discard most of the available ring structure as opposed to R_2^d . We remark that this issue hasn't been addressed by [LWW20]. It seems that for too narrow bounds on the entropic secret distribution, the number of available secrets is much smaller in the canonical embedding compared to the number with regard to the coefficient embedding.

Chapter 3

Classical Hardness of Module LWE

This chapter can be seen as a continuation of Chapter 2 and is therefore based on the same joint work with Corentin Jeudy, Adeline Roux-Langlois and Weiqiang Wen which is published in the proceedings of the conference Asiacrypt 2020 [BJRW20]. A presentation of 20 minutes has been recorded for the conference, illustrating the contributions and high level techniques of the paper.¹

Contents

3.1	Introduction	57
3.1.1	Our Contributions	58
3.1.2	Related Work	61
3.1.3	Roadmap	62
3.2	Classical Reduction for Exponentially Large Moduli	62
3.2.1	Step 1: From Gaussian Decoding Problem to M-LWE	62
3.2.2	Step 2: Classical Hardness for Exponentially Large Modulus	67
3.3	Binary Hardness and Adapted Error Distribution	69
3.4	Modulus Reduction	70
3.4.1	The General Case	71
3.4.2	The Case of Power-of-2 Cyclotomics	72

3.1 Introduction

In the following, we continue to study the module variant of LWE, denoted by M-LWE. The major advantage of LWE over other cryptographic hardness assumptions is its worst-case to average-case connection to well-studied problems over Euclidean lattices, for suitable parameter choices. In the seminal work of Regev [Reg05, Reg09], it is shown that any efficient solver for the average-case problem LWE can be transformed into an efficient solver for any instance of some worst-case lattice problem, such as finding a set of short independent vectors (SIVP) or the decision variant of finding short vectors (GapSVP). A standard relaxation of those two problems consists in solving them only up to an approximation factor γ , denoted by SIVP_γ and GapSVP_γ , respectively. One drawback of Regev’s reduction is that it is quantum, and thus an efficient solver for LWE only leads to an efficient *quantum* solver for SIVP_γ or GapSVP_γ , which needs to

¹<https://www.youtube.com/watch?v=zDdvV52FBMo>

operate on a quantum computer. A classical hardness proof for LWE starting from GapSVP_γ is shown by Peikert [Pei09], at the expense of requiring an exponentially large modulus q for the resulting LWE instance. However, such a modulus makes LWE rather impractical. Later, Brakerski et al. [BLP⁺13] use Peikert’s result and a modulus reduction for LWE to show a classical hardness proof for LWE while now allowing for a polynomially large modulus.

When Langlois and Stehlé [LS15] originally introduce and study the module variant of LWE, they prove an analogue of Regev’s reduction for M-LWE. In other words, they show for suitable parameters a worst-case to average-case quantum reduction from the SIVP_γ problem over module lattices to M-LWE. As for [Reg05], this reduction is quantum. In their introduction, Langlois and Stehlé [LS15] claim that Peikert’s dequantization [Pei09] carries over to the module case, making it plausible that a classical hardness proof should also be feasible for M-LWE. However, a formal proof for the classical hardness of M-LWE was left as an open problem.

3.1.1 Our Contributions

In this chapter, we make progress towards proving the classical hardness of M-LWE. Our main contribution is a classical reduction from the module variant of GapSVP_γ , which we denote by Mod-GapSVP_γ , with module rank at least 2 to M-LWE for any polynomial-sized modulus p and module rank d at least $3n + \omega(\log_2 n)$, where n is the degree of the underlying number field.

At a high level, we follow the structure of the classical hardness proof of LWE from Brakerski et al. [BLP⁺13]. To do so, we need three ingredients: First, a classical reduction for M-LWE with an exponential-sized modulus. As a second component, we need the hardness of M-LWE using a binary secret, and finally, a modulus reduction technique.

Classical Reduction: Regarding the first ingredient, we prove in Section 3.2 with Theorem 3.2 that Peikert’s dequantization [Pei09] carries over to the module case. The proof idea is the same as the one from Peikert, but with two novelties. First, we look at the structured variants of the corresponding problems, Mod-GapSVP_γ and M-LWE, where the underlying ring R is the ring of integers of a number field K . Recall that the problem Mod-GapSVP_γ becomes easy in the special case of modules of rank 1 (i.e., ideals) as the minimum of ideal lattices can be bounded above *and* below, see Lemma 1.4. This is why we restrict our results to modules of rank at least 2. Second, we replace the main component, a reduction from the Bounded Distance Decoding (BDD) problem to the search version of LWE, by the reduction from the Gaussian Decoding Problem (GDP) over modules to the decision version of M-LWE (Lemma 3.3, adapted from [PRS17a]). As a side contribution, we thus generalize the hardness of the decision variant of M-LWE to all number fields K , not only cyclotomic fields as in [LS15]. This is interesting, as we mainly use decision problems (and not their search variants) as the underlying hardness assumptions for cryptographic constructions.

Binary Secret: For the second ingredient, we can use our results from the previous chapter on the hardness of bin-M-LWE. More precisely, we make use of the improved reduction proven in Theorem 2.2. Recall that Corollary 2.1 is the instantiation of this theorem for power-of-2 cyclotomics. Note that in the original paper published at Asiacrypt 2020 [BJRW20], the improved reduction didn’t yet exist and we thus used the simple reduction of Theorem 2.1. However, this makes the overall classical hardness proof more involved, as we have to handle different transitions from the search to the decision variants and from discrete to continuous Gaussian error distributions. Hence, the reduction as shown in this thesis is not only better in parameters, but also requires less intermediate steps, as compared to the original paper [BJRW20].

Modulus Reduction: Finally, we provide a modulus reduction technique, the last required ingredient, where the rank of the underlying module is preserved. This corresponds to the modulus reduction for LWE shown by Brakerski et al. [BLP⁺13, Cor. 3.2]. Prior to this paper, Albrecht

and Deo [AD17a] adapted the more general result from [BLP⁺13, Thm. 3.1], from which the necessary Corollary 3.2 is deduced. Thus, in Section 3.4, we first recall their general result [AD17a, Thm. 1] and then derive Corollary 3.2, that we need for our purposes, from it. The quality of the latter depends on the underlying ring structure and how the binary secret distribution behaves. For the case of power-of-2 cyclotomics, we provide a concrete instantiation in Corollary 3.3. This involves the computation of upper bounds of the singular values of the rotation matrix. Note that Langlois and Stehlé [LS15] prove a modulus switching result for M-LWE from modulus q to modulus p , but the error increases at least by a multiplicative factor $\frac{q}{p}$, which is exponential if q is exponentially and p only polynomially large. Further, the reason why we need to go through the binary variant of M-LWE is because we want to keep the noise amplification during the modulus switching part as small as possible.

We now explain how to complete the proof of the classical hardness of decision M-LWE for any polynomial-sized modulus p and module rank d at least $3n + \omega(\log_2 n)$, as stated in Theorem 3.1. We further instantiate it for power-of-2 cyclotomics in Corollary 3.1. We refer to Figure 3.1 for an overview of the full proof.

In a first step, we classically reduce Mod-GapSVP_γ to M-LWE in Theorem 3.2, requiring the resulting modulus q to be exponentially large in the ring degree n , i.e., $q \geq 2^n$ (as $k \geq 2$). Further, the error distribution of M-LWE is given by Υ_α for some parameter α . This is a distribution over a particular set of elliptical Gaussian distributions in the canonical embedding, which we formally defined in Section 1.3.2. Then, we show in Lemma 3.4 how to move from Υ_α to the set of spherical Gaussians $\Psi_{\leq \alpha'}$, where $\alpha' = \alpha \cdot \omega(\log_2 n)$. In a next step, we use the binary secret result Theorem 2.2 of Chapter 2 to reduce M-LWE to bin-M-LWE, where the resulting rank d has to be at least $3 \cdot \log_2(q) + \omega(\log_2 n)$. As q is exponentially large in n , this leads to a rank d that is linear in n . Note that the prime number q has to fulfill some number-theoretic constraints which simplify to $q = 5 \pmod{8}$ in the case of power-of-2 cyclotomics (Corollary 2.1). Then, using Corollary 3.2, we show a reduction from bin-M-LWE with modulus q and Gaussian parameter bound β to bin-M-LWE with modulus p and Gaussian parameter bound β' , where $q \geq p \geq 1$ and $(\beta')^2 \geq (\sqrt{2}\beta)^2 + \Delta$. The parameter Δ is determined by the underlying ring R and is $\text{poly}(n)$ for the case of power-of-2 cyclotomics, as shown in Corollary 3.3. To conclude the classical hardness result of decision M-LWE with polynomial-sized modulus p , we trivially reduce bin-M-LWE to M-LWE by re-randomizing the secret.

The main contribution of this chapter is summarized in the following theorem.

Theorem 3.1 (Classical Hardness of M-LWE for Cyclotomic Fields)

Let $\nu \in \mathbb{N}$ with prime-power factorization $\nu = \prod_j p_j^{e_j}$. Further, let K be the ν -th cyclotomic field of degree $n = \varphi(\nu)$, and R its ring of integers. Set $\mu = \prod_j p_j$ and let q be a prime number such that $q = 1 \pmod{\mu}$, $\text{ord}_\nu(q) = \nu/\mu$ and $q > \max(2n, \mathfrak{s}_1(\mu)^{\varphi(\mu)})$, where $\mathfrak{s}_1(\mu)$ denotes the largest singular value of the Vandermonde matrix of the μ -th cyclotomic field. Let $d, p \in \mathbb{N}$ and $\beta', \gamma \in \mathbb{R}^+$. There is a classical PPT reduction from Mod-GapSVP_γ to $\text{M-LWE}_{n,d,p,\Psi_{\leq \beta'}}$, where $d \geq 3n + \omega(\log_2 n)$, $p = \text{poly}(n)$ and

$$\beta' = \tilde{\Theta} \left(\frac{n^{\frac{5}{2}}}{\gamma} \right).$$

In the special case of power-of-2 cyclotomic fields, the theorem implies the following corollary.

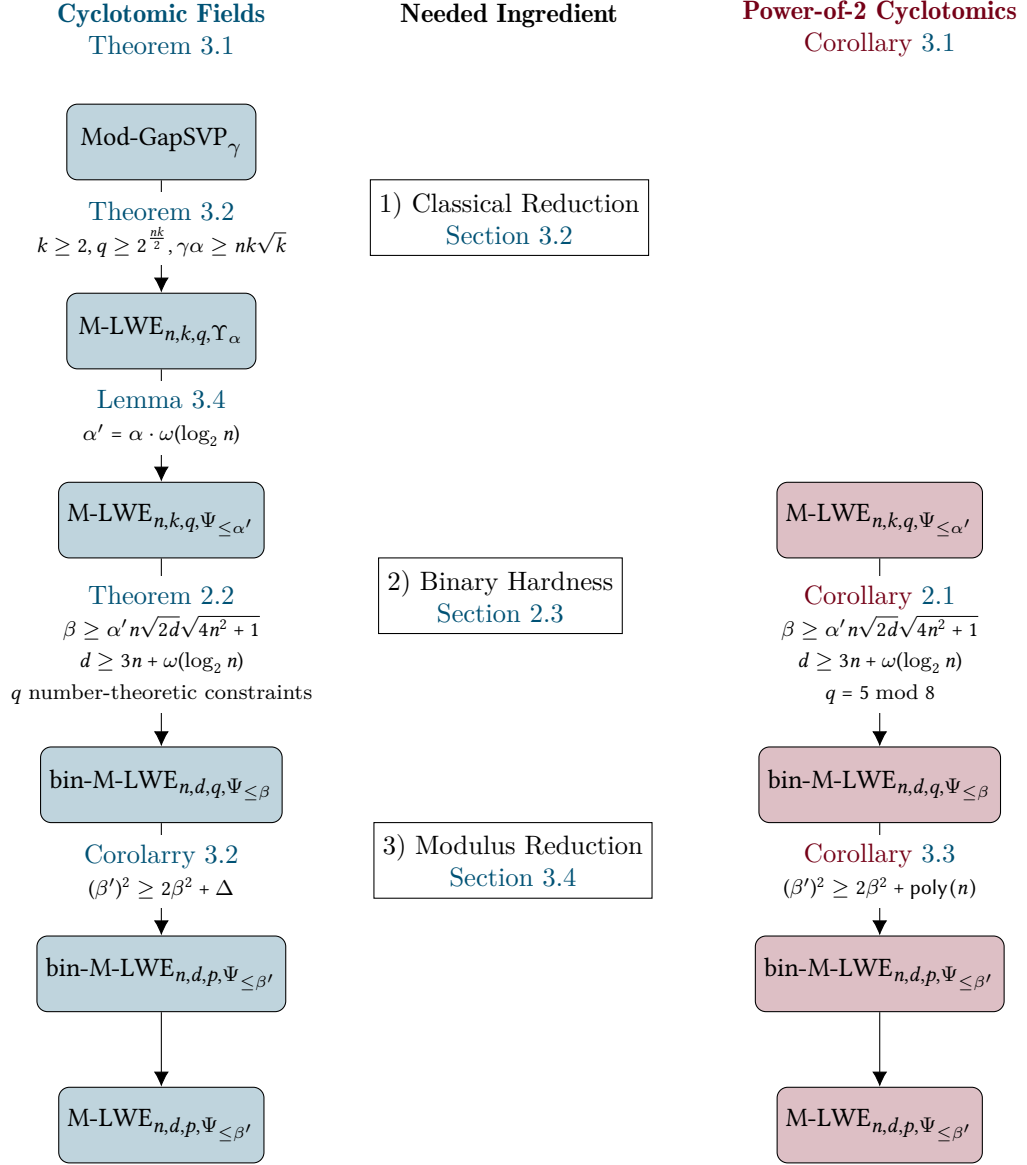


Figure 3.1: Overview of the complete classical hardness proof of decision M-LWE for linear rank d and polynomially large modulus p , as stated in Theorem 3.1 for K the ν -th cyclotomic number field of degree n . The right side shows the concrete instantiation for power-of-2 cyclotomics, as stated in Corollary 3.1. The parameter Δ is determined by the underlying ring R and is $\text{poly}(n)$ for the case of power-of-2 cyclotomics. For the reduction from M-LWE to its binary variant, we use the improved reduction from Chapter 2.

Corollary 3.1 (Classical Hardness of M-LWE for Power-of-2 Cyclotomics)

Let $\nu \in \mathbb{N}$ be a power of two, K be the ν -th cyclotomic field of degree $n = \nu/2$, and R its ring of integers. Let $q \in \mathbb{N}$ be prime such that $q \equiv 5 \pmod{8}$. Further, let $d, p \in \mathbb{N}$ and $\beta', \gamma \in \mathbb{R}^+$. There is a classical PPT reduction from Mod-GapSVP_γ to $\text{M-LWE}_{n,d,p,\Psi_{\leq \beta'}}$, where $d \geq 3n + \omega(\log_2 n)$ and

$$\beta' = \tilde{\Theta} \left(\frac{n^{\frac{5}{2}}}{\gamma} \right).$$

Using the advanced binary hardness proof from Section 2.3 instead of the simple reduction from Section 2.2 improves the β' from our original result [BJRW20, Thm. 2] by a factor of $\sqrt{m} \cdot n^{\frac{1}{4}}$, where m denotes the number of given samples. On the other hand, it increases the rank condition by an additive term of n .

3.1.2 Related Work

Classical Hardness of Ring LWE

As our overall classical reduction with polynomially large modulus is restricted to modules of rank linear in the ring degree, it doesn't apply to R-LWE, which is the special case of M-LWE where the rank d equals 1. Furthermore, the classical reduction with exponentially large modulus of Section 3.2 is meaningless for rank 1 modules, as Mod-GapSVP_γ is easy to solve in this case. However, as we discuss in the following, there are two different ways to obtain the classical hardness of R-LWE with an *exponentially* sized modulus.

The first is informally mentioned in [BLP⁺13] and can be achieved in two steps. First, by a dimension-modulus switching as in [BLP⁺13], we can reduce LWE in dimension d and with modulus q to LWE in dimension 1 and with modulus q^d with a slightly increased error rate. Then, by a ring switching technique as in [GHPS12], n samples of the latter one can be reduced to one sample of R-LWE over a ring of degree n and modulus q^d , while keeping the same error rate. For more details on the second step, we refer to [AD17a, App. B].

On the other hand, as a direct application of our classical hardness result of M-LWE, we can provide an alternative classical hardness result for R-LWE with exponential-sized modulus. The idea is that, using a rank-modulus switching as in [AD17a, WW19], we can instead reduce from M-LWE over d -rank modules of n -degree ring and modulus q , to R-LWE with n -degree ring and modulus q^d (and rank equals 1), 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 R-LWE over n -degree ring and q^d modulus where $d = O(n)$. Then, the underlying problem is the standard GapSVP over all Euclidean lattices of dimension $O(\sqrt{n})$ for the first result, while it is Mod-GapSVP over rank-2 module lattices of some $O(n)$ -degree ring for the second one.

Algebraically Structured LWE

A recent result from Peikert and Pepin [PP19] tightly proves the hardness of M-LWE over a number field K of degree n and with rank d assuming the hardness of R-LWE over any one of a class of number field extensions K'/K with extension degree $d = [K' : K]$. Instead of showing a modulus-rank trade-off as in [AD17a], they provide a degree-rank trade-off, where the underlying ring structure is changed, while preserving the modulus q . As pointed out by the authors themselves, this result is from a different nature than the original hardness result of M-LWE,

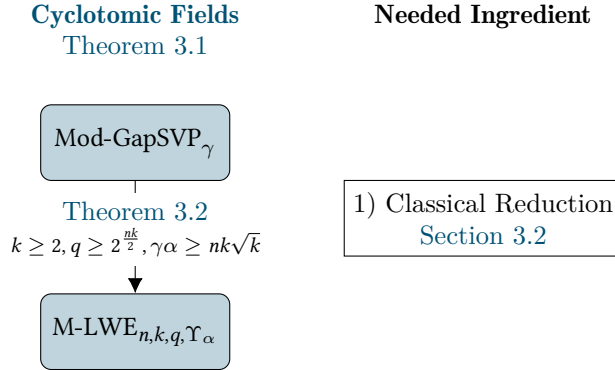


Figure 3.2: The first ingredient required for the complete classical hardness proof of decision M-LWE, as stated in Theorem 3.1 for K the ν -th cyclotomic number field of degree n .

given by Langlois and Stehlé [LS15]. The latter one established a reduction from worst-case problems over module lattices, whereas the reduction in [PP19] starts from an average-case problem over rings. The main intention of their paper is to unify the landscape of different algebraically structured variants of LWE.

3.1.3 Roadmap

The rest of the chapter is structured as follows. In Section 3.2, we prove a classical reduction for M-LWE with an exponentially large modulus. To this end, we first prove a reduction from the Gaussian Decoding Problem to M-LWE, which is afterwards used as a main building block for the classical proof. In Section 3.3, we explain how to use the binary hardness result from the previous chapter and how to adapt the error distribution. Finally, in Section 3.4 we provide a modulus-reduction first for the case of general number fields with general secret distributions and then for the special case of power-of-2 cyclotomics with secrets of small norm.

3.2 Classical Reduction for Exponentially Large Moduli

The goal of this section is to provide Theorem 3.2, which is the first ingredient needed to prove the classical hardness of M-LWE, as illustrated in Figure 3.2. To do so, we adapt the classical hardness reduction of LWE for exponentially large moduli from Peikert [Pei09, Thm. 3.1] to the module setting. In their introduction, Langlois and Stehlé [LS15] claim that Peikert’s dequantization [Pei09] carries over to the module case and we prove this claim in the following. By using the more recent results of Peikert et al. [PRS17a], our reduction directly reduces Mod-GapSVP to the decision variant M-LWE and holds for any number field K .

3.2.1 Step 1: From Gaussian Decoding Problem to M-LWE

As a first step, we prove a reduction from the Gaussian Decoding Problem (GDP) over module lattices (Definition 1.6) to M-LWE, which we then use later in Step 2 for the classical hardness proof. Informally speaking, GDP is a computational problem over lattices that, given a point \mathbf{y} in the Euclidean space with the promise that it is close to a given lattice, asks to recover the closest lattice point \mathbf{x} , or equivalently their difference $\mathbf{e} = \mathbf{y} - \mathbf{x}$. The particularity of GDP is

that the difference \mathbf{e} follows a Gaussian distribution. To obtain the required result, we prove the following Lemma 3.3, which is an adaptation of [PRS17a, Lem. 6.6] from ideals to modules. To do so, we essentially need to adapt two components from the original proof. The first is a transformation from an instance of the Bounded Distance Decoding (BDD) problem over some ideal lattice to samples of R-LWE. Fortunately, there is already in [LS15] an adaptation of it to the module setting which transforms an instance of BDD over some module lattice (Definition 1.4) to a sample of M-LWE using some Gaussians over the module lattice. Recall, that BDD is closely related to GDP, with the only difference that the vector \mathbf{e} doesn't have to follow a Gaussian distribution, but only needs to be of bounded norm. The second component which is important in the proof of [PRS17a, Lem. 6.6] is the Oracle Hidden Center Problem (OHCP, Definition 3.1) together with an efficient solver for it. This tool is independent of the underlying lattice and thus can still be used in the case of module lattices. As our proof is a direct adaptation of the original one, we strongly recommend to read the original paper with a more detailed introduction to the OHCP problem. We now recall the results for both tools.

Throughout this section, let K be a number field of degree n with R its ring of integers. As explained in Section 1.2.2, any R -module $\mathcal{M} \subseteq K^k$ of rank $k \geq 2$ can be identified with a module lattice of dimension $N = nk$. By \mathcal{M}^\vee we denote its dual lattice as defined in Section 1.1.5. Recall from Section 1.1.2 that any vector $\mathbf{e} \in K^k$ is equipped with the $(2, \infty)$ -norm defined by $\|\mathbf{e}\|_{2, \infty} = \max_{j \in [n]} (\sum_{\ell \in [k]} |\sigma_j(e_\ell)|^2)^{1/2}$. We further refer to Section 1.3.2 for the formal definitions of the different Gaussian distributions that arise in this section (as for instance the continuous Gaussian $D_{\mathbf{r}'}$, the discrete Gaussian distribution $\mathcal{D}_{\mathcal{M}, \mathbf{r}}$ and the distribution over elliptical Gaussians Υ_α). Further, recall the definition of the M-LWE distribution $A_{\mathbf{s}, \psi}^{\mathcal{M}}$ from Section 1.4.3.

Lemma 3.1 (Adapted from [LS15, Lem. 4.14])

Let K be a number field of degree n with ring of integers R and $k \in \mathbb{N}$. There exists a PPT algorithm that takes as input an integer $q \geq 2$ with known factorization, a module $\mathcal{M} \subseteq K^k$, a coset $\mathbf{e} + \mathcal{M}^\vee$ and bound $\delta \geq \|\mathbf{e}\|_{2, \infty}$, a parameter $r \geq \sqrt{2}q \cdot \eta_\varepsilon(\mathcal{M})$, and Gaussian samples from $\mathcal{D}_{\mathcal{M}, \mathbf{r}}$ for some $\mathbf{r} \geq r$. It outputs samples that are within negligible statistical distance of the M-LWE distribution $A_{\mathbf{s}, \psi}^{\mathcal{M}}$, for a uniformly random $\mathbf{s} \in (R_q^\vee)^k$ and $\psi = D_{\mathbf{r}'}$, where the coefficients of \mathbf{r}' are given for every $j \in [n]$ by

$$(r'_j)^2 = \left(\frac{r\delta}{q}\right)^2 + \frac{r_j^2 \sum_{\ell \in [k]} |\sigma_j(e_\ell)|^2}{q^2}.$$

We now formally define the Oracle Hidden Center Problem. Note that we use the wording of the problem as in the updated version from June 2020 of [PRS17a], accessible via the IACR eprint server [PRS17b]. The modifications from this updated version have, however, no impact on our proof of Lemma 3.3.

Definition 3.1 (Oracle Hidden Center Problem): Let $k \in \mathbb{N}$, $d \in \mathbb{R}^+$, $\beta \geq 1$ and $\varepsilon, \delta \in [0, 1) \cap \mathbb{R}$, the $(\varepsilon, \delta, \beta)$ -OHCP is an approximate search problem defined as follows. Given access to an oracle $\mathcal{O}: \mathbb{R}^k \times \mathbb{R}^+ \rightarrow \{0, 1\}$, such that its acceptance probability $p(\mathbf{z}, t)$ for any input (\mathbf{z}, t) with $\|\mathbf{z} - \mathbf{z}^*\|_2 \leq \beta d$ depends only on $\exp(t) \cdot \|\mathbf{z} - \mathbf{z}^*\|_2$ for some hidden center $\mathbf{z}^* \in \mathbb{R}^k$ with $\delta d \leq \|\mathbf{z}^*\|_2 \leq d$, the goal is to output $\mathbf{z} \in \mathbb{R}^k$ such that $\|\mathbf{z} - \mathbf{z}^*\|_2 \leq \varepsilon \cdot d$.

Lemma 3.2 (Solver for OHCP [PRS17a, Prop. 4.4])

There is a PPT algorithm that takes as input a confidence parameter $\kappa \geq 20 \log_2(k+1)$, the scale parameter d , and solves $(\exp(-\kappa), \exp(-\kappa), 1 + 1/\kappa)$ -OHCP in dimension k except with probability $\exp(-\kappa)$, provided that exists a $p(\infty) \in [0, 1]$ such that

1. $p(\mathbf{0}, t^*) - p(\infty) \geq 1/\kappa$ for some $t^* \in \mathbb{R}^+$;
2. $|p(\mathbf{0}, t) - p(\infty)| \leq 2 \exp(-t/\kappa)$ for any $t \in \mathbb{R}^+$; and
3. $p(\mathbf{z}, t)$ is κ -Lipschitz in t for any $\mathbf{z} \in \mathbb{R}^k$.

Let $n = t_1 + 2t_2$ the decomposition of the degree n into the number of real and complex embeddings. As in [PRS17a], for any $r > 0, \zeta > 0$, and $T \geq 1$, we define the set of non-spherical parameter vectors $W_{r,\zeta,T}$ as the set of cardinality $(t_1 + t_2) \cdot (T + 1)$, containing for each $\ell \in [t_1 + t_2]$ and $j \in \{0, \dots, T\}$ the vector $\mathbf{r}_{\ell j}$ which is equal to r in all coordinates except in the ℓ -th (and the $(\ell + t_2)$ -th if $\ell > t_1$), where it is equal to $r \cdot (1 + \zeta)^j$. We are now ready to state and prove the reduction from GDP to M-LWE.

Lemma 3.3 (Adapted from [PRS17a, Lem. 6.6])

Let K be a number field of degree n with ring of integers R and $k \in \mathbb{N}$. There exists a PPT algorithm that, given an oracle that solves $\text{M-LWE}_{q,T,\alpha}$, for a real $\alpha \in (0, 1)$ and an integer $q \geq 2$ together with its factorization, a rank k module $\mathcal{M} \subseteq K^k$, a parameter $r \geq \sqrt{2q} \cdot \eta_\varepsilon(\mathcal{M})$ for $\varepsilon = \exp(-kn)$, and polynomially many samples from the discrete Gaussian distribution $\mathcal{D}_{\mathcal{M},\mathbf{r}}$ for each $\mathbf{r} \in W_{r,\zeta,T}$ (for some $\zeta = 1/\text{poly}(n)$ and $T = \text{poly}(n)$), solves $\text{GDP}_{\mathcal{M}^\vee, g}$, for $g = \alpha q / (\sqrt{2kr})$.

Note that if we sample a coset $\mathbf{e} + \mathcal{M}^\vee$ as in $\text{GDP}_{\mathcal{M}^\vee, g}$ with $g = \alpha q / \sqrt{2kr}$ its $(2, \infty)$ -norm is bounded above by $\delta = \sqrt{k} \cdot g \cdot G(n) = \alpha q G(n) / \sqrt{2r}$, for some fixed $G(n) = \omega(\sqrt{\log_2 n})$. If we then apply the lemma above with δ and $r_j = r$, the distribution of the resulting error rate \mathbf{r}' is exactly Υ_α .

Overall, it suffices to carefully adapt the definition of the BDD instances to get M-LWE instances. Then we can use the oracle for M-LWE to simulate some oracles \mathcal{O}_j such that \mathcal{O}_j has hidden center $\sigma_j(\mathbf{e})$ for every $j \in [n]$. An approximation of $\sigma_j(\mathbf{e})$ can then be used with Babai's algorithm to reconstruct \mathbf{e} . Still, the proof maintains to be quite technical.

Proof: *Assumption on α :* Fix q, α, k, r, g as in the statement of the lemma. By $G(n)$ we denote a function in $\omega(\sqrt{\log_2 n})$. Recall that any R -module $\mathcal{M} \subseteq K^k$ of rank k corresponds to a module lattice of dimension $N = nk$. Assume $\alpha \leq \exp(-N)$. Using Lemma 1.8 with $\varepsilon = \exp(-N)$, it yields with high probability

$$\|\sigma(\mathbf{e})\|_2 \leq \sqrt{N}g \leq \frac{\alpha q \sqrt{N}}{\sqrt{2kr}} \leq \frac{\alpha}{2\sqrt{k}} \cdot \frac{\sqrt{N}}{\eta_\varepsilon(\mathcal{M})} \leq \frac{\alpha \sqrt{\pi}}{2\sqrt{k}} \cdot \lambda_1(\mathcal{M}^\vee) \leq 2^{-N} \lambda_1(\mathcal{M}^\vee).$$

Thus, the LLL algorithm (Lemma 1.5) would solve the GDP instance efficiently and without loss of generality we may assume in the following that $\alpha > \exp(-N)$.

The high level idea of the reduction: Define $\kappa = \text{poly}(N)$ with $\kappa \geq 100N^2 m$, such that the

advantage of the oracle for $\text{M-LWE}_{q, \Gamma_\alpha}$ is at least $2/\kappa$, when it receives m input samples. Let $\mathbf{e} + \mathcal{M}^\vee$ denote the input to the $\text{GDP}_{\mathcal{M}^\vee, g}$ problem, where $\mathbf{e} = (e_1, \dots, e_k)$ with $e_\ell \leftarrow D_g$ for every $\ell \in [k]$. The goal is to recover \mathbf{e} .

The reduction uses the M-LWE oracle to simulate $t_1 + t_2$ different oracles \mathcal{O}_j , such that the acceptance probability of \mathcal{O}_j on input (\mathbf{z}, t) only depends on $\exp(t) \cdot \|\mathbf{z} - \sigma_j(\mathbf{e})\|_2$, where σ denotes the canonical embedding. In other words, \mathcal{O}_j has a hidden center $\sigma_j(\mathbf{e})$, defining an instance of the OHCP as in Definition 3.1. More concretely, we define for the t_1 real embeddings the oracles $\mathcal{O}_j: \mathbb{R}^k \times \mathbb{R}^+ \rightarrow \{0, 1\}$ (for $j \in [t_1]$) and for the t_2 complex embeddings the oracles $\mathcal{O}_j: \mathbb{C}^k \times \mathbb{R}^+ \rightarrow \{0, 1\}$ (for $t_1 + 1 \leq j \leq t_1 + t_2$). The reduction can thus use the efficient solver of Lemma 3.2 to find a good approximation of $\sigma_j(\mathbf{e})$ for every $j \in [t_1 + t_2]$. Note that $\sigma_{t_1+t_2+j}(\mathbf{e})$ is the complex conjugate of $\sigma_{t_1+j}(\mathbf{e})$, so it suffices to have t_2 oracles for the complex embeddings.

The oracles: In order to define the oracle \mathcal{O}_j we define the following functions. For the real embeddings with $j \in [t_1]$, we set $k_j: \mathbb{R}^k \rightarrow (K_{\mathbb{R}})^k$ with $k_j(\mathbf{z}) = \sigma^{-1}(z_1 \cdot \mathbf{e}_j, \dots, z_k \cdot \mathbf{e}_j)^T$. For the complex embeddings with $t_1 + 1 \leq j \leq t_1 + t_2$, we set $k_j: \mathbb{C}^k \rightarrow (K_{\mathbb{R}})^k$ with $k_j(\mathbf{z}) = \sigma^{-1}(z_1 \cdot \mathbf{e}_j + \bar{z}_1 \cdot \mathbf{e}_{j+t_2}, \dots, z_k \cdot \mathbf{e}_j + \bar{z}_k \cdot \mathbf{e}_{j+t_2})^T$. Recall that \bar{z} denotes the complex conjugate of the complex number $z \in \mathbb{C}$ and that \mathbf{e}_j denotes the j -th unit vector. On input (\mathbf{z}, t) , the oracle \mathcal{O}_j chooses samples from $\mathcal{D}_{\mathcal{M}, r, i_j}$, where the index i is defined by t as $(1 + \zeta)^i = \exp(t)$. It then runs the transformation from Lemma 3.1 on these Gaussian samples, the coset $\mathbf{e} - k_j(\mathbf{z}) + \mathcal{M}^\vee$, parameter r and distribution bound $\delta = \frac{\alpha q G(n)}{\sqrt{2r}} = \sqrt{k} \cdot g \cdot G(n)$. Note that $\|\mathbf{e} - k_j(\mathbf{z})\|_{2, \infty} \leq \|\mathbf{e}\|_{2, \infty} + \|\mathbf{z}\|_2 \leq \delta + \|\mathbf{z}\|_2$ and that we mostly only care about the behavior of the oracle when $\mathbf{z} = 0$. Let $A_{j, \mathbf{z}, t}^{\mathcal{M}}$ be the resulting M-LWE samples. The oracle \mathcal{O}_j then calls the oracle for M-LWE on these samples and finally outputs 1 if the latter one accepts. In a next step, the reduction uses the efficient solver for the OHCP (Lemma 3.2) with confidence parameter κ , distance bound δ and receives some approximation \mathbf{z}_j to the oracle's hidden center $\sigma_j(\mathbf{e})$. Finally, it runs the LLL algorithm on the coset $\mathbf{e} - \sigma^{-1}(\mathbf{z}_1, \dots, \mathbf{z}_n)^T + \mathcal{M}^\vee$, receives as output $\hat{\mathbf{e}}$ and returns $\hat{\mathbf{e}} + \sigma^{-1}(\mathbf{z}_1, \dots, \mathbf{z}_n)^T$ as solution to the GDP instance.

Proof of the claims: We now prove that this reduction works as claimed. First, we assume that the \mathbf{z}_j are valid solutions to the $(\exp(-\kappa), \exp(-\kappa), 1 + 1/\kappa)$ -OHCP with hidden center $\sigma_j(\mathbf{e})$ for every $j \in [t_1 + t_2]$ and show that the reduction outputs the correct answer. Since \mathbf{z}_j is a valid solution, by definition and Lemma 1.8 we have

$$\|\sigma_j(\mathbf{e}) - \mathbf{z}_j\|_2 \leq \exp(-\kappa)\delta = \exp(-\kappa) \frac{\alpha q G(n)}{\sqrt{2r}} \leq \exp(-\kappa) \frac{G(n)}{\eta_\varepsilon(\mathcal{M})} \leq 2^{-N} \frac{\lambda_1(\mathcal{M}^\vee)}{\sqrt{N}}.$$

Having $\sigma_j(\mathbf{e})$ for all $j \in [t_1 + t_2]$, we construct the full vector $\sigma(\mathbf{e})$ and compute $\|\sigma(\mathbf{e}) - (\mathbf{z}_1, \dots, \mathbf{z}_n)^T\|_2 \leq (\sum_{j \in [n]} \|\sigma_j(\mathbf{e}) - \mathbf{z}_j\|_2^2)^{1/2} \leq 2^{-N} \lambda_1(\mathcal{M}^\vee)$, and finally $\|\mathbf{e} - \sigma^{-1}(\mathbf{z}_1, \dots, \mathbf{z}_n)^T\|_2 \leq 2^{-N} \lambda_1(\mathcal{M}^\vee)$. Thus, the LLL algorithm (Lemma 1.5) succeeds and outputs the correct element $\hat{\mathbf{e}} = \mathbf{e} - \sigma^{-1}(\mathbf{z}_1, \dots, \mathbf{z}_n)^T$ and thus $\hat{\mathbf{e}} + \sigma^{-1}(\mathbf{z}_1, \dots, \mathbf{z}_n)^T$ is a correct solution to the GDP instance.

Second, we show that for every $j \in [t_1 + t_2]$ the oracle \mathcal{O}_j represents a valid instance of $(\exp(-\kappa), \exp(-\kappa), 1 + 1/\kappa)$ -OHCP with hidden center $\sigma_j(\mathbf{e})$. Lemma 3.1 with $r_j = r \exp(t)$ implies that the distribution of $A_{j, \mathbf{z}, t}^{\mathcal{M}}$ only depends on $\exp(t) \cdot \|\mathbf{z} - \sigma(\mathbf{e})\|_2$. As $\kappa = \text{poly}(N)$ and every coefficient e_ℓ of \mathbf{e} is sampled from D_g , it yields for $\delta = \sqrt{k} \cdot g \cdot G(n)$ that $\exp(-\kappa)\delta \leq \|\sigma_j(\mathbf{e})\|_2 \leq \delta$, for every $j \in [t_1 + t_2]$ except with negligible probability. Thus, \mathcal{O}_j, κ and δ correspond indeed to a valid instance of OHCP.

Efficient solver for OHCP: To conclude the proof, we show that for every $j \in [t_1 + t_2]$ the

oracle \mathcal{O}_j satisfies the requirements to apply Lemma 3.2, which provides an efficient solver for OHCP. Let $p_j(\mathbf{z}, t)$ denote the probability that \mathcal{O}_j outputs 1 on input (\mathbf{z}, t) and $p(\infty)$ the probability that the M-LWE oracle outputs 1 on uniformly random input. Further, it yields $p_j(\mathbf{0}, 0) = p_\ell(\mathbf{0}, 0)$ for all $j, \ell \in [t_1 + t_2]$. Also note that $p_j(\mathbf{0}, 0) - p(\infty)$ corresponds to the advantage of the M-LWE oracle, where the error rate that we obtain from Lemma 3.1 on input $\mathbf{e} + \mathcal{M}^\vee$ is exactly Υ_α (as $k_j(\mathbf{z}) = \mathbf{0}$ and $\exp(t) = 1$). As we assume that this advantage is bounded below by $2/\kappa$, it yields $\mathbb{E}_{\mathbf{e} \leftarrow D_g} [p_j(\mathbf{0}, 0) - p(\infty)] \geq 2/\kappa$.

For the first item of Lemma 3.2, we have to show that $p_j(\mathbf{0}, t^*) - p(\infty) \geq 1/\kappa$ for some $t^* \in \mathbb{R}^+$. Setting $t^* = 0$ and using a Markov argument, we get that $\Pr_{\mathbf{e} \leftarrow D_g} [p_j(\mathbf{0}, 0) - p(\infty) \geq 1/\kappa] \geq 1/(\kappa - 1)$, which is non-negligible. For the second item of Lemma 3.2, we have to show that $|p_j(\mathbf{0}, t) - p(\infty)| \leq 2 \exp(-t/\kappa)$ for any $t \in \mathbb{R}^+$. It suffices to show that the distribution of $A_{j, \mathbf{0}, t}^{\mathcal{M}}$ is within statistical distance $2 \exp(-t/\kappa)$ of the uniform distribution for any $t \in \mathbb{R}^+$. Let $t \geq \ln(2)\kappa$. Using [PRS17a, Lem. 6.9], the statistical distance for m samples is bounded above by

$$\begin{aligned} m \cdot \varepsilon &= m \cdot \exp(-c^2 n) = m \cdot \exp(-n \prod_{\ell \in [n]} (r'_\ell)^{2/n}) \\ &\leq m \cdot \exp\left(-n/q^2 \prod_{\ell \in [n]} (r_\ell)^{2/n} \prod_{\ell \in [n]} (\|\sigma_\ell(\mathbf{e})\|_2)^{2/n}\right) \\ &= m \cdot \exp\left(-n/q^2 \exp(2t/n) r^2 \prod_{\ell \in [n]} (\|\sigma_\ell(\mathbf{e})\|_2)^{2/n}\right). \end{aligned}$$

Here, we use that $(r'_\ell)^2 \geq ((r_\ell)^2 \|\sigma_\ell(\mathbf{e})\|_2^2)/q^2$ for all $\ell \in [n]$ and $r_\ell = r$ for all $\ell \neq j$ and $r_j = r \exp(t)$. We deduce for every $\ell \in [n]$ that

$$\|\sigma_\ell(\mathbf{e})\|_2 \geq \exp(-N) \cdot \delta = \exp(-N) \cdot \frac{\alpha q G(n)}{\sqrt{2r}} > \exp(-2N - 1) q/r.$$

We plug this equation into the one above to obtain

$$\begin{aligned} m \cdot \varepsilon &\leq m \cdot \exp\left(-n/q^2 \exp(2t/n) r^2 \prod_{\ell \in [n]} (\|\sigma_\ell(\mathbf{e})\|_2)^{2/n}\right) \\ &\leq m \cdot \exp(-n \exp(2t/n - 4N - 2)) \\ &< m \cdot \exp((-n(t/\kappa n + \log_2(m/2)/n)) = 2 \exp(-t/\kappa), \end{aligned}$$

where we used in the third step the loose equation $\exp(2t/n - 4N - 2) > \exp(t/n) > t/\kappa n + \log_2(m/2)/n$, as we required $t \geq \ln(2)\kappa \geq \kappa/10 \geq 10N^2 m$.

And finally, for the last item of Lemma 3.2, we have to show the Lipschitz property, i.e., that $|p_j(\mathbf{z}, t_1) - p_j(\mathbf{z}, t_2)| \leq \kappa |t_1 - t_2|$ for any $t_1, t_2 \in \mathbb{R}^+$. As mentioned before, the distribution of $A_{j, \mathbf{z}, t}^{\mathcal{M}}$ only depends on $\exp(t) \|\mathbf{z} - \sigma_j(\mathbf{e})\|_2$. Thus, by Lemma 1.19, the distributions of $A_{j, \mathbf{z}, t_1}^{\mathcal{M}}$ and $A_{j, \mathbf{z}, t_2}^{\mathcal{M}}$ are within statistical distance $\min\{1, 10m(\exp(|t_1 - t_2|) - 1)\} \leq \kappa |t_1 - t_2|$. Here, we use that if $|t_1 - t_2| > 1/\kappa$, then the bound $\kappa |t_1 - t_2|$ is trivial. On the other hand, if $|t_1 - t_2| \leq 1/\kappa \ll 1/100$, then $\exp(|t_1 - t_2|) - 1 \ll 10|t_1 - t_2|$, so $10m(\exp(|t_1 - t_2|) - 1) \leq \kappa |t_1 - t_2|$, which completes the proof. \blacksquare

3.2.2 Step 2: Classical Hardness for Exponentially Large Modulus

Using the results from above, we are able to adapt the classical hardness result of LWE from Peikert [Pei09, Thm. 3.1] to modules, yielding for suitable parameters a reduction from GapSVP over module lattices in the worst-case to the M-LWE problem.

Theorem 3.2 (Classical Hardness of M-LWE)

Let $\alpha, \gamma \in \mathbb{R}^+$ such that $\alpha \in (0, 1)$ and let $n, k, q \in \mathbb{N}$, defining $N = nk$. We consider a number field K of degree n with R its ring of integers. Further, assume that $k \geq 2$, $q \geq 2^{\frac{N}{2}}$ and $\gamma \geq \frac{N\sqrt{k}}{\alpha}$. Let $\mathcal{M} \subseteq K^k$ be a rank- k module. There exists a PPT reduction from solving Mod-GapSVP_γ in the worst-case to solving the problem $\text{M-LWE}_{n,k,q,\gamma,\alpha}$, using $\text{poly}(N)$ samples.

The proof idea is the same as the one from Peikert, but with two modifications. First, we look at the module variants of the corresponding problems, i.e., Mod-GapSVP over module lattices and M-LWE, where the underlying ring R is the ring of integers of a number field K . Second, we replace the main component, a reduction from the BDD problem to the search version of LWE ([Pei09, Prop. 3.4], originally from [Reg05, Lem. 3.4]), by the reduction proven above from the GDP problem over modules to the decision version of M-LWE (Lemma 3.3). As Lemma 3.3 requires access to polynomially many samples from some discrete Gaussian distribution over a lattice, we make use of the Gaussian sampler from Lemma 1.12.

Proof: Let $\mathcal{M} \subseteq K^k$ be a rank- k module over R , such that the corresponding module lattice of dimension N has basis $\mathbf{B} = (\mathbf{b}_j)_{j \in [N]}$. Further, let δ be a positive real. The Mod-GapSVP_γ problem as stated in Definition 1.2 asks to decide whether $\lambda_1(\mathcal{M}) \leq \delta$ (YES instance) or $\lambda_1(\mathcal{M}) > \gamma\delta$ (NO instance). Without loss of generality, we assume that the basis \mathbf{B} is LLL-reduced (Lemma 1.5) and appropriately scaled, thus the following three conditions hold:

- C1) $\lambda_1(\mathcal{M}) \leq 2^{\frac{N}{2}}$,
- C2) $\min_{j \in [N]} \|\text{GS}(\mathbf{b}_j)\|_2 \geq 1$,
- C3) $1 \leq \delta \leq 2^{\frac{N}{2}}/\gamma$.

Recall that we denote by $\text{GS}(\mathbf{B}) = (\text{GS}(\mathbf{b}_j))_{j \in [N]}$ the Gram-Schmidt orthogonalization of \mathbf{B} from left to right. Note for C3, that Mod-GapSVP_γ becomes trivial if δ lies outside this range: If $\delta < 1$, using C2, we know that $\lambda_1(\mathcal{M}) \geq 1$ and thus it is definitely not a YES instance. If $\delta > 2^{\frac{N}{2}}/\gamma$, using C1, we know that $\delta\gamma > 2^{\frac{N}{2}} \geq \lambda_1(\mathcal{M})$ and thus it is definitely not a NO instance.

The reduction: The reduction executes the following procedure $\text{poly}(N)$ many times:

- Choose $\mathbf{w} \leftarrow D_{g'}$ with $g' = \frac{\delta}{2} \cdot \sqrt{N}$,
- Compute $\mathbf{w} + \mathcal{M}$,
- Run the GDP_g oracle from Lemma 3.3 with $\mathbf{w} + \mathcal{M}$, $r = \frac{q\sqrt{2N}}{\gamma\delta}$, $g = \frac{\alpha q}{\sqrt{2kr}}$, and using the Gaussian sampler from Lemma 1.12,

- Compare the output of the oracle with \mathbf{w} .

If the oracle's answer is always correct, output NO, otherwise YES.

Gaussian sampler: First, we show that the Gaussian sampler from Lemma 1.12 always succeeds to provide polynomially many samples from the discrete Gaussian distribution $\mathcal{D}_{\mathcal{M}^\vee, \mathbf{r}}$ for each $\mathbf{r} \in W_{r, \zeta, T}$ (for some $\zeta = 1/\text{poly}(n)$ and $T = \text{poly}(n)$), needed in Lemma 3.3. Note that for every $\mathbf{r} = (r_j)_{j \in [n]} \in W_{r, \zeta, T}$ it yields $r_j \geq r$ for every $j \in [n]$. As the task of sampling Gaussians becomes easier with increasing parameter, it suffices to show that the Gaussian sampler succeeds for r . By Section 1.2.1, $\mathbf{D} = (\mathbf{B}^{-1})^T$ defines a basis of the dual \mathcal{M}^\vee , where we denote by \mathbf{d}_j its column vectors for $j \in [N]$. It yields for the Euclidean norm that $\|\text{GS}(\mathbf{D})\|_2 = \|\text{GS}(\mathbf{B})\|_2^{-1}$. As we require in condition C2 that $\min_{j \in [N]} \|\text{GS}(\mathbf{b}_j)\|_2 \geq 1$, it follows $\max_{j \in [N]} \|\text{GS}(\mathbf{d}_j)\|_2 \leq 1$. Using the condition C3 and that $q \geq 2^{\frac{N}{2}}$, it yields

$$r = \frac{q\sqrt{2N}}{\gamma\delta} \geq \sqrt{2N} \geq 1 \cdot \sqrt{\ln(2N+4)/\pi} \geq \max_{j \in [N]} \|\text{GS}(\mathbf{d}_j)\|_2 \cdot \sqrt{\ln(2N+4)/\pi},$$

and thus the Gaussian sampler from Lemma 1.12 always succeeds.

Case 1 (NO instance): Now, we assume that the reduction is given a NO instance, i.e., $\lambda_1(\mathcal{M}) > \gamma\delta$. We claim that in this case, all requirements from Lemma 3.3 are fulfilled and thus the oracle always successfully decodes the GDP instance and hence always outputs the correct answer. Using Lemma 1.8 it yields $\eta_\varepsilon(\mathcal{M}^\vee) \leq \sqrt{N}/\lambda_1(\mathcal{M})$ for $\varepsilon = \exp(-N)$. Thus,

$$r = \frac{q\sqrt{2N}}{\gamma\delta} > \frac{q\sqrt{2N}}{\lambda_1(\mathcal{M})} \geq \sqrt{2}q \cdot \eta_\varepsilon(\mathcal{M}^\vee).$$

Further, \mathbf{w} is sampled from $D_{g'}$ with

$$g' = \frac{\delta}{2} \cdot \sqrt{N} \leq \frac{\alpha\gamma\delta}{2\sqrt{nk}} = \frac{\alpha q}{\sqrt{2kr}} = g,$$

where we use the definitions of g and r and the condition on $\alpha\gamma$ from the theorem statement, i.e., $\alpha\gamma \geq N\sqrt{k}$. Additionally, \mathbf{w} is the unique solution to this problem as with high probability

$$2 \cdot \|\mathbf{w}\|_2 \leq 2 \cdot g' \sqrt{nk} = 2 \cdot \frac{\delta}{2} \cdot \sqrt{N} \cdot \sqrt{nk} \leq \frac{\alpha\gamma\delta}{\sqrt{k}} < \gamma\delta < \lambda_1(\mathcal{M}).$$

Thus, the difference $\mathbf{w} - \mathbf{w}'$ of two different solutions $\mathbf{w} \neq \mathbf{w}'$ to the problem would lead to a non-zero lattice point in \mathcal{M} of norm smaller than the minimum, which is a contradiction of the definition of the minimum $\lambda_1(\mathcal{M})$.

Case 2 (YES instance): If, on the other hand, the reduction is given a YES instance, i.e., $\lambda_1(\mathcal{M}) \leq \delta$, we can consider the following alternate experiment. Let \mathbf{z} be a shortest vector in \mathcal{M} with $\|\mathbf{z}\|_2 = \lambda_1(\mathcal{M}) \leq \delta$. Now, we replace \mathbf{w} by $\mathbf{w}' = \mathbf{w} + \mathbf{z}$ in the second step of the reduction and thus hand in $\mathbf{w}' + \mathcal{M}$ to the GDP oracle. Using the probability preservation property of the statistical distance (Lemma 1.17) of \mathbf{w} and \mathbf{w}' , it yields

$$\begin{aligned} \Pr[\mathcal{R}(\mathbf{w} + \mathcal{M}) = \mathbf{w}] &\leq \Delta(\mathbf{w}; \mathbf{w}') + \Pr[\mathcal{R}(\mathbf{w}' + \mathcal{M}) = \mathbf{w}'] \\ &\leq \Delta(\mathbf{w}; \mathbf{w}') + 1 - \Pr[\mathcal{R}(\mathbf{w}' + \mathcal{M}) = \mathbf{w}], \end{aligned}$$

where \mathcal{R} denotes the GDP oracle. Note that $\mathbf{w}' + \mathcal{M} = \mathbf{w} + \mathcal{M}$, so in the real experiment we

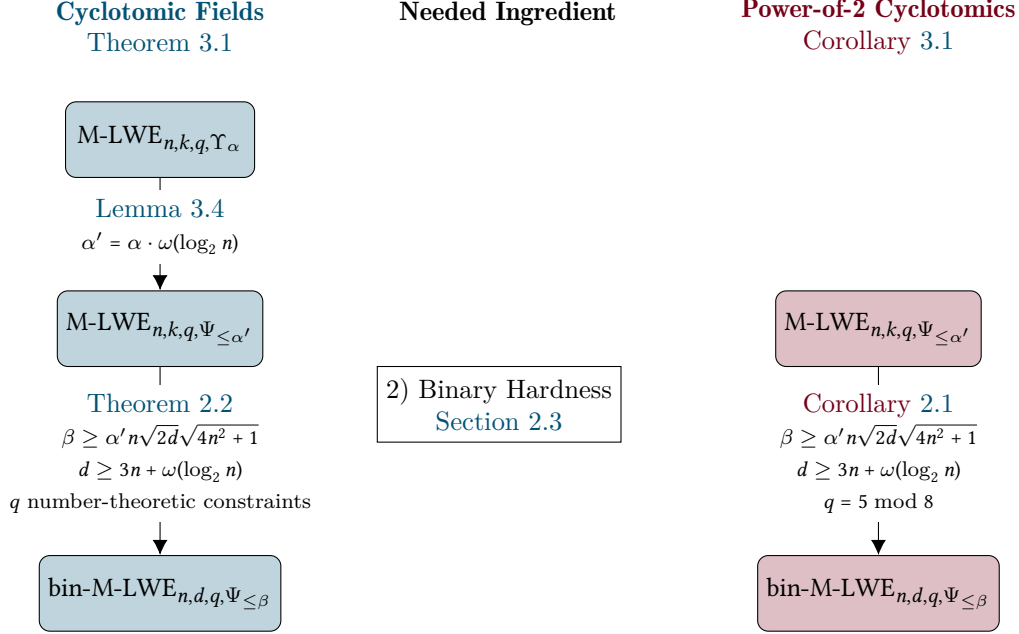


Figure 3.3: The second ingredient required for the complete classical hardness proof of decision M-LWE, as stated in Theorem 3.1 for K the ν -th cyclotomic number field of degree n . The right side shows the concrete instantiation for power-of-2 cyclotomics, as stated in Corollary 3.1. For the reduction from M-LWE to its binary variant, we use the improved reduction from Chapter 2.

have $\Pr[\mathcal{R}(\mathbf{w}' + \mathcal{M}) = \mathbf{w}] = \Pr[\mathcal{R}(\mathbf{w} + \mathcal{M}) = \mathbf{w}]$ and thus

$$\Pr[\mathcal{R}(\mathbf{w} + \mathcal{M}) = \mathbf{w}] \leq \frac{1 + \Delta(\mathbf{w}; \mathbf{w}')}{2}.$$

Using the statistical distance of two Gaussian distributions with the same width but different means, Lemma 1.18, we obtain

$$\Delta(\mathbf{w}; \mathbf{w}') \leq \frac{\sqrt{2\pi} \|\mathbf{z}\|_2}{g'} \leq \frac{\sqrt{2\pi} \delta}{g'} \leq 2 \frac{\sqrt{2\pi}}{\sqrt{N}},$$

where we use that $g' = \delta \sqrt{N}/2$ and thus $\Pr[\mathcal{R}(\mathbf{w} + \mathcal{M}) = \mathbf{w}] \leq \frac{1}{2} + \frac{\sqrt{2\pi}}{\sqrt{N}}$. For sufficiently many iterations, the oracle gives a wrong answer in at least one iteration and the reduction outputs YES. ■

3.3 Binary Hardness and Adapted Error Distribution

As the second ingredient for our classical hardness result for M-LWE, as depicted in Figure 3.3, we need the binary hardness of M-LWE. Fortunately, Chapter 2 provides a detailed study of this problem. Concretely, we use Theorem 2.2 from Section 2.3 as it directly proves the hardness of

the decision version of bin-M-LWE. The instantiation for power-of-2 cyclotomics is provided in Corollary 2.1. To complete the proof of Theorem 3.1, we need to connect the first and second ingredient by adapting the error distribution after the classical reduction and before the binary secret proof. More precisely, we have to move from the distribution Υ_α over elliptical Gaussian distributions, as used within Section 3.2 and defined in Section 1.3.2, to the set of continuous Gaussian distributions $\Psi_{\leq \alpha'}$, as used in Section 2.3. To achieve this, we use similar techniques as in [LS15, Sec. 4.4]. Note that their definition of Υ is slightly different than ours, positively influencing the bound we obtain.

Lemma 3.4

Let n, k, q be positive integers and α be a positive real. There exists a PPT reduction from $\text{M-LWE}_{n,k,q,\Upsilon_\alpha}$ to $\text{M-LWE}_{n,k,q,\Psi_{\leq \alpha'}}$, where $\alpha' = \alpha \cdot \omega(\log_2 n)$.

Proof: Our goal is to reduce $\text{M-LWE}_{n,k,q,\Upsilon_\alpha}$ to $\text{M-LWE}_{n,k,q,\Psi_{\leq \alpha'}}$, where α' is given by $\alpha \cdot \omega(\log_2 n)$. Let K be a number field of degree n and let $n = t_1 + 2t_2$ be the decomposition of n in real and complex embeddings. Recall, that Υ_α is a distribution over elliptical Gaussian distributions $D_{\mathbf{r}}$, where r_j^2 is distributed as a shifted chi-squared distribution for the real embeddings (i.e., $j \in [t_1]$) and as a shifted chi-squared distribution with two degrees of freedom for complex embeddings (i.e., $j \in \{t_1 + 1, \dots, t_1 + t_2\}$). Using properties about chi-squared distributions (see for instance [LM00, Lem. 1]), it yields that $r_j \leq \frac{\alpha}{\sqrt{2}} \cdot \omega(\log_2 n) \leq \alpha \cdot \omega(\log_2 n) = \alpha'$ with probability negligible close to 1. Thus, $\text{M-LWE}_{n,k,q,\Psi_{\leq \alpha'}}$ is not easier than $\text{M-LWE}_{n,k,q,\Upsilon_\alpha}$. ■

3.4 Modulus Reduction

Recall from Figure 3.1 that we need three ingredients to prove the classical hardness of M-LWE. In Section 3.2, we provided the first ingredient. For the second ingredient, the hardness of M-LWE with binary secrets, we can use our contributions from Chapter 2. The goal of this section is to present the last ingredient required to complete the classical hardness proof. To this end, we show a modulus reduction for bin-M-LWE, where the rank of the underlying module is preserved, as illustrated in Figure 3.4. This corresponds to the modulus reduction for LWE shown by Brakerski et al. [BLP⁺13, Cor. 3.2]. Note that Langlois and Stehlé [LS15] prove a modulus switching result from $\text{M-LWE}_{n,d,q,\Upsilon_\beta}$ to $\text{M-LWE}_{n,d,p,\Upsilon_{\beta'}}$, where the error increases at least by a multiplicative factor $\frac{q}{p}$. If q is exponential-sized and p only polynomial-sized, this factor is exponentially large and thus not suitable for our intentions.

Prior to this thesis, Albrecht and Deo [AD17a] adapt the more general result from Brakerski et al. [BLP⁺13, Thm. 3.1], from which Corollary 3.2 is deduced. Thus, we first recall their general result [AD17a, Thm. 1] and then derive the corollary we need from it. We use their more recent and simplified version as updated on the IACR eprint server [AD17b, Thm. 1]. The theorem provides a transformation that maps M-LWE samples for the modulus q and rank d with a Gaussian noise distribution of parameter β to M-LWE samples for a different modulus p , rank d' and Gaussian noise distribution of parameter β' . The secret \mathbf{s} of the first sample is related to the secret of the second sample via some matrix \mathbf{G} . Simultaneously, it maps the uniform distribution to the uniform distribution over the corresponding sets.




Figure 3.4: The third (and last) ingredient required for the complete classical hardness proof of decision M-LWE, as stated in Theorem 3.1 for K the ν -th cyclotomic number field of degree n . The right side shows the concrete instantiation for power-of-2 cyclotomics, as stated in Corollary 3.1. The parameter Δ is determined by the underlying ring R and is $\text{poly}(n)$ for the case of power-of-2 cyclotomics.

Theorem 3.3 (Modulus Switching [AD17b, Thm. 1])

Let K be a number field of degree n with R its ring of integers. Let $d, d', q, p \in \mathbb{N}$, $\varepsilon \in (0, \frac{1}{2})$ and $\mathbf{G} \in R^{d' \times d}$. Fix a vector $\mathbf{s} = (s_1, \dots, s_d)^T \in (R_q^\vee)^d$. Further, let Λ be the lattice given by $\Lambda = \frac{1}{p} \mathbf{G}_H^T R^{d'} + R^d$ with known basis \mathbf{B}_Λ in the canonical embedding, let \mathbf{B}_R be some known basis of R in H and let r be a real number such that

$$r \geq \max \left(\|\text{GS}(\mathbf{B}_\Lambda)\|_2, \frac{1}{q} \|\text{GS}(\mathbf{B}_R)\|_2 \right) \cdot \sqrt{2 \ln(2nd(1 + 1/\varepsilon))/\pi},$$

where $\text{GS}(\cdot)$ denotes the Gram-Schmidt orthogonalization from left to right. There exists an efficient mapping $\mathcal{F}: (R_q)^d \times \mathbb{T}_{R^\vee} \rightarrow (R_p)^{d'} \times \mathbb{T}_{R^\vee}$ such that:

1. The output distribution $\mathcal{F}(U((R_q)^d \times \mathbb{T}_{R^\vee}))$ given uniform input is within statistical distance 4ε of the uniform distribution over $(R_p)^{d'} \times \mathbb{T}_{R^\vee}$.
2. The output distribution $\mathcal{F}(A_{\mathbf{s}, D_\beta}^M)$ is within statistical distance $(2d+6)\varepsilon$ of $A_{\mathbf{G}_s, D_{\beta'}}^M$, where

$$(\beta'_k)^2 = \beta^2 + r^2(\gamma^2 + \sum_{j \in [d]} |\sigma_k(s_j)|^2),$$

for $k \in [n]$ and γ satisfying $\gamma^2 \geq B^2 d$ with $B = \|\mathbf{s}\|_\infty$.

3.4.1 The General Case

Whereas Albrecht and Deo [AD17a] derive from the theorem above a rank-modulus trade-off, defining a map from M-LWE with module rank d and modulus q to M-LWE with rank d/k and modulus q^k for any divisor k of d , we are interested in another particular instance of Theorem 3.3 where the rank of the module is preserved. The following corollary specializes the general result to the case of $\mathbf{G} = \mathbf{I}_d \in R^{d \times d}$ and its proof is very similar to the one of [AD17a]. Overall, we

obtain a modulus reduction, where the rank d is preserved.

Corollary 3.2 (Modulus Reduction for General Number Fields)

Let $d, q, p \in \mathbb{N}$ with $q \geq p$ and $\varepsilon, \beta \in \mathbb{R}^+$ with $\varepsilon \in (0, \frac{1}{2})$ and $\mathbf{G} = \mathbf{I}_d \in R^{d \times d}$. Let ψ be a distribution over R_q^\vee satisfying

$$\Pr_{s \leftarrow \psi} \left[\max_{k \in [n]} |\sigma_k(s)| > B \right] \leq \delta$$

for some non-negative real numbers B and δ . By ψ^d we denote the distribution over $(R_q^\vee)^d$, where every coefficient is sampled from ψ independently. Let \mathbf{B}_R be some known basis of R in H and r be a real number such that

$$r \geq \frac{1}{p} \|\text{GS}(\mathbf{B}_R)\|_2 \cdot \sqrt{2 \ln(2nd(1 + 1/\varepsilon))/\pi}.$$

Then, there is a polynomial-time reduction from M-LWE $_{n,d,q,\Psi_{\leq \beta}}^m$ with secret distribution ψ^d to the problem M-LWE $_{n,d,p,\Psi_{\leq \beta'}}^m$ with secret distribution ψ^d for $(\beta')^2 \geq \beta^2 + 2r^2 B^2 d$. This reduction reduces the advantage by at most $1 - (1 - \delta)^d + (2d + 6)\varepsilon m$.

Proof: We use the transformation from Theorem 3.3 by taking $\gamma^2 = B^2 d$ and replacing $\sum_{j \in [d]} |\sigma_k(s_j)|^2$ for every $k \in [n]$ by $B^2 d$. We can write \mathbf{G} in the coefficient embedding as $\hat{\mathbf{G}} = \mathbf{I}_d \otimes \mathbf{I}_n = \mathbf{I}_{dn}$, where \otimes denotes the Kronecker product of two matrices. This defines the corresponding lattice $\hat{\Lambda} = \frac{1}{p} \hat{\mathbf{G}}^T \mathbb{Z}^{dn} + \mathbb{Z}^{dn}$ with basis $\mathbf{B}_{\hat{\Lambda}} = \frac{1}{p} \mathbf{I}_{dn}$. To move from the coefficient embedding to the canonical embedding, we can simply multiply the basis by the matrix $\mathbf{B}_{R^d} = \mathbf{I}_d \otimes \mathbf{B}_R$. The basis for $\Lambda = \frac{1}{p} \mathbf{G}_H^T R^d + R^d$ given in the canonical embedding is thus given by

$$\mathbf{B}_\Lambda = \left(\frac{1}{p} \mathbf{I}_d \otimes \mathbf{I}_n \right) \cdot \left(\mathbf{I}_d \otimes \mathbf{B}_R \right) = \frac{1}{p} \mathbf{I}_d \otimes \mathbf{B}_R,$$

using the mixed product property of the Kronecker product. Orthogonalizing from left to right gives $\|\text{GS}(\mathbf{B}_\Lambda)\|_2 = \frac{1}{p} \|\text{GS}(\mathbf{B}_R)\|_2$. As $q \geq p$, we have $\frac{1}{q} \|\text{GS}(\mathbf{B}_R)\|_2 \leq \frac{1}{p} \|\text{GS}(\mathbf{B}_R)\|_2 = \|\text{GS}(\mathbf{B}_\Lambda)\|_2$ and thus r satisfies the condition of Theorem 3.3. The loss in advantage is the result of a simple probability calculus. The event that $\max_{k \in [n]} |\sigma_k(s)| \leq B$ happens with probability greater than $1 - \delta$. As the secret vector $\mathbf{s} = (s_1, \dots, s_d)^T \in (R_q^\vee)^d$ is chosen by drawing d times independently from ψ , we have to add the advantage loss of $1 - (1 - \delta)^d$ to the one coming from Theorem 3.3. ■

3.4.2 The Case of Power-of-2 Cyclotomics

The quality of Corollary 3.2 depends on the term $\Delta = 2r^2 B^2 d$, that we have to add to the error width β^2 in order to obtain the resulting error parameter $(\beta')^2$. This factor is determined by the rank d , the first bound B on the secret distribution ψ and the number r , which itself is quantified by the field degree n , the starting modulus q , the reduced modulus p and the norm $\|\text{GS}(\mathbf{B}_R)\|_2$.

In the following, we give a concrete calculation example for those parameters in the case of power-of-2 cyclotomic rings for the secret distribution ψ given by the uniform distribution

over R_2^\vee . Let ν be a power of two, defining the ring of integers of the ν -th cyclotomic field, given by $R = \mathbb{Z}[\zeta] \cong \mathbb{Z}[x]/\langle f(x) \rangle$, where $f(x) = x^\nu + 1$ and $n = \nu/2$.

Corollary 3.3 (Modulus Reduction for Power-of-2 Cyclotomics and Binary Secrets)

Let R be a cyclotomic ring of degree n , where n is a power of 2. Let $d, q, p \in \mathbb{N}$ with $q \geq p$ and $\varepsilon, \beta \in \mathbb{R}^+$ with $\varepsilon \in (0, \frac{1}{2})$ and $\mathbf{G} = \mathbf{I}_d \in R^{d \times d}$. Let further be r a real number such that

$$r \geq \frac{1}{p} \sqrt{n} \cdot \sqrt{2 \ln(2nd(1 + 1/\varepsilon))/\pi}.$$

For $(\beta')^2 \geq \beta^2 + 2dr^2$, there is a polynomial-time reduction from the problem M-LWE $_{n,d,q,\Psi \leq \beta}^m$ with secret distribution $U((R_2^\vee)^d)$ to M-LWE $_{n,d,p,\Psi \leq \beta'}^m$ with the same secret distribution $U((R_2^\vee)^d)$. This reduction reduces the advantage by at most $(2d + 6)\varepsilon m$.

In order to guarantee a negligible loss in advantage, we require m and d to be polynomial in the security parameter and ε negligibly small. For p polynomial in n and fulfilling $p \geq \frac{\sqrt{2dn} \cdot \sqrt{2 \ln(2nd(1+1/\varepsilon))/\pi}}{\beta}$, we can make sure that $r = \frac{\beta}{\sqrt{2d}}$ is a valid choice, implying that $\beta' \geq \sqrt{2}\beta$. Hence, the noise is only increased by a constant multiplicative factor.

Proof: To prove the claim it suffices to show a first condition on the norm of \mathbf{B}_R , i.e., that $\|\text{GS}(\mathbf{B}_R)\|_2 = \sqrt{n}$ and as a second condition that the uniform distribution over R_2^\vee satisfies the requirement of Corollary 3.2 for the parameters $\delta = 0$ and $B = 1$.

First Condition: Let R be a cyclotomic ring of degree n , where n is a power of 2. As explained in Example 1.3, its dual R^\vee is just a scaling of the ring R itself, given by $R^\vee = \frac{1}{n}R$. Further, following Example 1.2, the Vandermonde matrix that maps the vector of an element in R defined by its canonical embedding to the vector corresponding to the coefficient embedding is a scaled isometry with scaling factor $\frac{1}{\sqrt{n}}$. A basis \mathbf{B}_R for R in H is given by $\sqrt{n} \cdot \mathbf{U}$, where \mathbf{U} is unitary and thus, $\|\text{GS}(\mathbf{B}_R)\|_2 = \sqrt{n}$.

Second Condition: For any element $s \in R$, let \mathbf{S}_H be the matrix of multiplication by s in the canonical embedding written in the basis $\{\mathbf{h}_j\}_{j \in [n]}$ of H . Let $\text{Rot}(s)$ be the matrix of multiplication by s in the coefficient embedding. As mentioned above, going from the coefficient embedding to the canonical embedding is a scaled isometry of scaling factor \sqrt{n} . Thus,

$$\mathbf{S}_H = (\mathbf{B}_R)^{-1} \cdot \text{Rot}(s) \cdot \mathbf{B}_R = \frac{1}{\sqrt{n}} \cdot \mathbf{U}^\dagger \cdot \text{Rot}(s) \cdot \sqrt{n} \cdot \mathbf{U} = \mathbf{U}^\dagger \cdot \text{Rot}(s) \cdot \mathbf{U},$$

where \mathbf{U} is unitary. As explained in Section 1.1.2, the singular values of \mathbf{S}_H are given by $|\sigma_j(s)|$ for $j \in [n]$. It yields

$$\begin{aligned} (\mathbf{S}_H)^\dagger \mathbf{S}_H &= (\mathbf{U}^\dagger \cdot \text{Rot}(s) \cdot \mathbf{U})^\dagger (\mathbf{U}^\dagger \cdot \text{Rot}(s) \cdot \mathbf{U}) \\ &= \mathbf{U}^{-1} \cdot \text{Rot}(s)^T \cdot \text{Rot}(s) \cdot \mathbf{U}. \end{aligned}$$

As a conclusion, the singular values of $\text{Rot}(s)$ are exactly the same as the one of \mathbf{S}_H , given by $|\sigma_j(s)|$ for $j \in [n]$. The largest singular value of $\text{Rot}(s)$ thus determines the maximum of the set $\{|\sigma_j(s)|\}_{j \in [n]}$. We use this observation to compute the bound B of Corollary 3.2 for the case where ψ equals $U((R_2^\vee)^d)$. Note that we provide new bounds, as the ones calculated

by Albrecht and Deo [AD17a] hold for a Gaussian, and not a binary secret distribution.

Using the identity $R_2^\vee = \frac{1}{n}R_2$, we can write $\text{Rot}(s) = \frac{1}{n}\text{Rot}(\tilde{s})$, where $\tilde{s} \in R_2$ and $\text{Rot}(\tilde{s})$ only has entries from the set $\{0, 1\}$. Here we use the special form of the rotation matrix for power-of-2 cyclotomics as explained in Example 1.2. Let $\text{Rot}(\tilde{s}) = \mathbf{U} \cdot \mathbf{\Sigma} \cdot \mathbf{V}^\dagger$ be the singular value decomposition of $\text{Rot}(\tilde{s})$, where \mathbf{U} and \mathbf{V} are unitary matrices over \mathbb{R} and $\mathbf{\Sigma}$ is a diagonal matrix with the singular values of $\text{Rot}(\tilde{s})$ on its diagonal. The singular value decomposition of $\text{Rot}(s)$ is thus given by $\text{Rot}(s) = \mathbf{U} \cdot \frac{1}{n}\mathbf{\Sigma} \cdot \mathbf{V}^\dagger$ and we can deduce that the singular values of $\text{Rot}(s)$ are just the singular values of $\text{Rot}(\tilde{s})$, shrank by a factor of $\frac{1}{n}$.

The largest singular value $\mathfrak{s}_1(\text{Rot}(\tilde{s}))$ of $\text{Rot}(\tilde{s})$ is bounded above by its Frobenius norm $\|\text{Rot}(\tilde{s})\|_F$ and hence

$$\mathfrak{s}_1(\text{Rot}(\tilde{s})) \leq \|\text{Rot}(\tilde{s})\|_F = \left(\sum_{i,j \in [n]} |\text{Rot}(\tilde{s})_{ij}|^2 \right)^{1/2} \leq n.$$

It follows $\mathfrak{s}_1(\text{Rot}(s)) \leq 1$. We can thus set $B = 1$ with $\delta = 0$. ■

Generalization to Larger Secrets

One may wonder if it is possible to generalize this modulus reduction to larger secret distributions, as we did for both binary hardness proofs of Chapter 2, see Section 2.4. More concretely, we consider the case where the secret's coefficients are not necessarily chosen uniformly at random over $\{0, 1\}$, but over a slightly larger set $\{0, \dots, \eta - 1\}$ for some positive integer $\eta \ll q$. As the only thing to adapt in the proof of Corollary 3.3 is the second condition, defining the bounds B and δ , we can positively answer this question, on the expense of a larger noise increase.

In the case where ψ equals the uniform distribution over $(R_\eta^\vee)^d$ (instead of $(R_2^\vee)^d$ as before) the same reasoning holds with respect to the largest singular value of the rotation matrix, leading to $B = \eta - 1$ and $\delta = 0$ and implying the following corollary.

Corollary 3.4 (Modulus Reduction for Power-of-2 Cyclotomics and Small Secrets)

Let R be a cyclotomic ring of degree n , where n is a power of 2. Let $d, q, p, \eta \in \mathbb{N}$ with $q \geq p$ and $\varepsilon, \beta \in \mathbb{R}^+$ with $\varepsilon \in (0, \frac{1}{2})$ and $\mathbf{G} = \mathbf{I}_d \in R^{d \times d}$. Let further be r a real number such that

$$r \geq \frac{1}{p} \sqrt{n} \cdot \sqrt{2 \ln(2nd(1 + 1/\varepsilon))/\pi}.$$

For $(\beta')^2 \geq \beta^2 + 2dr^2(\eta - 1)^2$, there is a polynomial-time reduction from the problem $\text{M-LWE}_{n,d,q,\Psi \leq \beta}^m$ with secret distribution $U((R_\eta^\vee)^d)$ to $\text{M-LWE}_{n,d,p,\Psi \leq \beta'}^m$ with the same secret distribution $U((R_\eta^\vee)^d)$. This reduction reduces the advantage by at most $(2d + 6)\varepsilon m$.

For p polynomial in n and fulfilling $p \geq \sqrt{2dn} \cdot (\eta - 1) \cdot \sqrt{2 \ln(2nd(1 + 1/\varepsilon))/\pi}/\beta$, we can make sure that $r = \beta/(\sqrt{2d}(\eta - 1))$ is a valid choice, implying that $\beta' \geq \sqrt{2}\beta$.

Chapter 4

Middle-Product Learning With Rounding

The content of this chapter is based on a joint work with Shi Bai, Dipayan Das, Adeline Roux-Langlois, Weiqiang Wen and Zhenfei Zhang which is published in the proceedings of the conference Asiacrypt 2019 [BBD⁺19].

Contents

4.1	Introduction	75
4.1.1	Our Contributions	76
4.1.2	Related Work	77
4.1.3	Roadmap	77
4.2	Random Hankel Matrices	78
4.3	Middle-Product Learning With Rounding (MP-LWR)	81
4.4	Hardness of Computational MP-LWR	82

4.1 Introduction

In Section 1.4, we defined the Middle-Product Learning With Errors (MP-LWE, Def. 1.15) problem, as introduced by Roşca et al. [RSSS17]. Recall that the hardness of MP-LWE is guaranteed as long as the Polynomial Learning With Errors (P-LWE, Def. 1.10) problem defined over some polynomial $f(x)$, belonging to an exponentially large family of polynomials, is hard. This makes the hardness of MP-LWE independent of a concrete choice of such a defining polynomial.

Additionally, we defined a deterministic variant of LWE, the Learning With Rounding (LWR, Def. 1.17) problem, as introduced by Banerjee et al. [BPR12]. It only requires sampling elements uniformly at random over finite sets, and doesn't need to sample from discrete Gaussian distributions. Gaussian distributions play a crucial role in the worst-case to average-case reduction of LWE [Reg05] and P-LWE [SSTX09] and are in general costly, difficult to implement and vulnerable to side-channel attacks, e.g., [DB15, BHLY16, Pes16, Saa18]. As for LWE, the LWR variant possesses a structured analogue over rings, called Ring Learning With Rounding (R-LWR).

In 2016, Bogdanov et al. [BGM⁺16] use the Rényi divergence to show a sample preserving reduction from the search variant of R-LWE to the search variant of R-LWR. In another work, Alperin-Sheriff and Apon [AA16] further improve the parameter sets for the reduction. In

particular, the reduction is dimension-preserving with a polynomial-sized modulus. However, a reduction from decision R-LWE to decision R-LWR with a polynomial-sized modulus was longtime an open problem, see Section 4.1.2 on related work for more details.

Nevertheless, due to the simplicity and efficiency of R-LWR, several encryption schemes as SABER [DKRV18] (third round) and Round5 [BBF⁺19] (second round) participating in the NIST standardization process are basing their hardness on (decision) R-LWR. To overcome the lack of provable hardness for decision R-LWR with practical parameters, Chen et al. [CZZ18] propose a new assumption, called the Computational Learning With Rounding Over Rings (R-CLWR) problem. They show a reduction from decision R-LWE to R-CLWR, where the secret in the R-LWE sample is drawn uniformly at random from the set of all invertible ring elements whose coefficients are small. They also show that one can construct an efficient PKE scheme based on the hardness of R-CLWR in the random oracle model.

4.1.1 Our Contributions

The main motivation of this work is to combine both variants MP-LWE and R-LWR in order to define a new problem that benefits from both of their advantages. To this end, we introduce a new hardness assumption which we refer to as the Middle Product Computational Learning With Rounding (MP-CLWR) problem. On the one hand, MP-CLWR uses rounding in a similar way to R-LWR and hence avoids the Gaussian error sampling. On the other hand, the hardness of MP-CLWR does not depend on a specific defining polynomial. Thus, the MP-CLWR assumption enjoys the security advantage of MP-LWE and the simplicity advantage of LWR.

The reason why we introduce the problem in its *computational* form, instead of the more standard *search* or *decision* variant, is twofold. On the one hand, it is as for today unclear how to reduce the hardness of decision MP-LWR from worst-case lattice problems, while maintaining the coefficient-wise rounding and allowing for a polynomially large modulus. On the other hand, it is unclear how to construct encryption schemes directly on the search variant. The computational variant gives a solution to this dilemma, as we can derive an efficient reduction still allowing for a polynomially large modulus and at the same time we can use it to build an encryption scheme (which we introduce separately in Chapter 6).

In the following, we give a brief overview of the MP-CLWR problem and our proof for its hardness. An MP-CLWR sample is given by $(a, b = \lfloor a \odot_d s \rfloor_p)$, where a is sampled from the uniform distribution over $\mathbb{Z}_q^{<n}[x]$ and s is a fixed element in $\mathbb{Z}_q^{<n+d-1}[x]$. Recall that $\mathbb{Z}_q^{<n}[x]$ denotes the set of polynomials in $\mathbb{Z}_q[x]$ with degree less than n and that $a \odot_d s$ denotes the middle-product of length d of the polynomials a and s , as introduced in Section 1.4.4. Further, note that $\lfloor \cdot \rfloor_p$ denotes the modular rounding function from Section 1.4.5. We define the MP-CLWR problem as the following game, where we embed 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 $(a_j, \lfloor a_j \odot_d s \rfloor_p)_{j \in [t]}$ and the source \mathcal{S}_2 of the second experiment provides t rounded uniform samples $(a_j, \lfloor b_j \rfloor_p)_{j \in [t]}$, where all a_j and b_j are independently sampled from the corresponding uniform distribution. The challenger \mathcal{C} now uses these samples to compute an **Input** and a **Target**. They send the **Input** to the adversary \mathcal{A} who themselves compute 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 in both experiments the same. The MP-CLWR assumption captures that an adversary has no more advantage to compute the correct output if they receive an input derived from rounded middle-product samples (Experiment 1) than if they get an input derived from rounded uniform samples (Experiment 2). We informally illustrate the experiment setting in Figure 4.1.

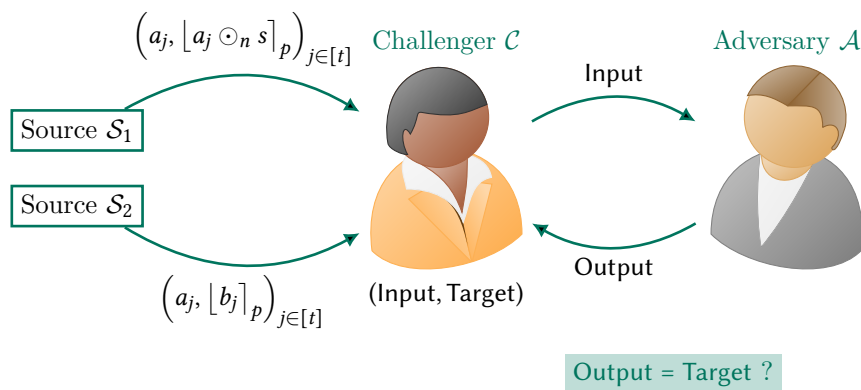


Figure 4.1: The experiment setting of the computational assumption (informal)

In order to show in Theorem 4.1 that the hardness of MP-CLWR is guaranteed by the hardness of MP-LWE, we take an instance (a, b) of the latter and round the second component b . This leads us to an instance of the so-called *rounded* MP-LWE problem (see Section 4.3). We can then use the Rényi divergence to measure the distance between instances of this rounded MP-LWE problem and instances from MP-LWR. It is easy to go from the decision version to the computational one, as the distinguishing game can be seen as a special case of the computational game (when setting the target as the bit to decide which distribution is used). The Hankel matrix plays an important role in the sequence of reductions as it can be used to represent the middle-product of two polynomials and thus it helps to analyze when the middle-product of a uniform random element is again uniform random.

4.1.2 Related Work

As stated before, a reduction from decision R-LWE to decision R-LWR with polynomial-sized modulus was longtime an open problem, motivating the introduction of intermediate problems such as the computational variant of R-LWR and similarly our newly defined MP-CLWR problem. After our results have been published, Liu and Wang [LW20] addressed this open problem by providing a search-to-decision reduction for R-LWR. In order to do so, they need to define a new way of rounding. Whereas in our work (and in all other works before) we round an element of the ring $\mathbb{Z}_q[x]/\langle f(x) \rangle$ coefficient-wise, they now round elements with respect to so-called *normal integral bases*. They further provide a reduction from R-LWR to MP-LWR for the search and decision variants. Putting both results together, we obtain the hardness of decision MP-LWR. However, the authors didn't provide an analysis how the different rounding notions are related to each other. This is why we think that it would be interesting to further study the relation of the different ways of rounding in future works.

4.1.3 Roadmap

The rest of the chapter is structured as follows. In Section 4.2, we first prove some new results on random Hankel matrices which we need in the hardness reduction and might be of independent interest. We give a formal definition of the MP-CLWR problem in Section 4.3 and show that it is at least as hard as decision MP-LWE in Section 4.4.

4.2 Random Hankel Matrices

In this section, we show new results on the distribution of random Hankel matrices. First, we recall the definition of Hankel and Toeplitz matrices for a given polynomial, which we interpret as usual as a vector. We prove a lower bound for the probability that the Hankel matrix of a polynomial which is chosen uniformly at random has full rank. Finally, this result leads to a uniformity property of the middle-product which plays a crucial part in the hardness reduction of the new middle-product learning with rounding assumption in Section 4.4.

Hankel and Toeplitz matrices are not only used in the context of the middle-product of two polynomials. More generally, as pointed out by Kaltofen and Lobo [KL96], Toeplitz matrices are used as pre-conditioners in the process of solving linear systems of equations having unstructured coefficient matrices. The attractiveness of these structured matrices is twofold: First, it suffices to store the first column and first row, in order to rebuild the whole matrix. Second, the product of a Toeplitz matrix and a vector is in fact a convolution and can be computed in quasi-linear time using the Fast Fourier Transformation.

Other than that, large-dimensional random matrices with additional algebraic structure, as Hankel and Toeplitz matrices, play an important role in statistics, in particular in multivariate analysis. More concretely, Hankel matrices arise in polynomial regressions and Toeplitz matrices appear as covariance of stationary processes. In particular, the spectral distribution for their eigenvalues is important and was studied by Bryc et al. [BDJ06].

We start by recalling the definitions of Hankel and Toeplitz matrices.

Definition 4.1 (Hankel Matrix): Let $q, n, d \in \mathbb{N}$ such that $d \leq n$ and $a \in \mathbb{Z}_q^{\leq n+d-1}[x]$ with coefficient vector $\mathbf{a} = (a_0, \dots, a_{n+d-2})^T$. We define the Hankel matrix of a of order $d + n - 1$ as

$$\text{Hank}(a) = \begin{bmatrix} a_0 & a_1 & \dots & a_{d-1} & \dots & a_{n-1} \\ a_1 & a_2 & \dots & a_d & \dots & a_n \\ \vdots & \vdots & \ddots & \vdots & & \vdots \\ a_{d-1} & a_d & \dots & a_{2d-2} & \dots & a_{n+d-2} \end{bmatrix} \in \mathbb{Z}_q^{d \times n}.$$

The Hankel matrix is fully determined by its first row and its last column. Its rank is at most d . If it has full rank d we write $\text{rank}(\text{Hank}(a)) = d$.

Definition 4.2 (Toeplitz Matrix): Let $q, n, d \in \mathbb{N}$ such that $d \leq n$ and $a \in \mathbb{Z}_q^{\leq n+d-1}[x]$ with coefficient vector $\mathbf{a} = (a_0, \dots, a_{n+d-2})^T$. The Toeplitz matrix of a of order $d + n - 1$ is given by

$$\text{Toep}(a) = \begin{bmatrix} a_0 & a_1 & \dots & a_{d-1} & \dots & a_{n-1} \\ a_n & a_0 & \dots & a_{d-2} & \dots & a_{n-2} \\ \vdots & \vdots & \ddots & \vdots & & \vdots \\ a_{n+d-2} & a_{n+d-3} & \dots & a_0 & \dots & a_{n-d} \end{bmatrix} \in \mathbb{Z}_q^{d \times n}.$$

The Toeplitz matrix is fully determined by its first row and its first column. There exists a special relation between Toeplitz matrices and the Hankel matrices. Let \mathbf{J}_n be the reflection matrix of order n defined as

$$\mathbf{J}_n = \begin{bmatrix} 0 & \cdots & 0 & 0 & 1 \\ 0 & \cdots & 0 & 1 & 0 \\ 0 & \cdots & 1 & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ 1 & \cdots & 0 & 0 & 0 \end{bmatrix} \in \mathbb{Z}_q^{n \times n}.$$

Then, for any polynomial $a \in \mathbb{Z}_q^{\leq n+d-1}[x]$ with coefficient vector $\mathbf{a} = (\mathbf{a}', \mathbf{a}'')$ in $\mathbb{Z}_q^n \times \mathbb{Z}_q^{d-1}$ it yields $\text{Toep}(a) \cdot \mathbf{J}_n = \text{Hank}(\tilde{a})$, where \tilde{a} is the polynomial given by the coefficient vector $\tilde{\mathbf{a}} = (\text{rev}(\mathbf{a}'), \mathbf{a}'')$ with $\text{rev}(\mathbf{a}')$ denoting the vector \mathbf{a}' in reverse order. Thus, we can use the result of Kaltofen and Lobo [KL96] about random Toeplitz matrices to calculate the probability of a random Hankel matrix to have full rank.

Lemma 4.1 (Rank of Random Hankel Matrices)

Let $q \in \mathbb{N}$ with unique prime power factorization given by $q = \prod_{j \in [\ell]} p_j^{\alpha_j}$, where p_j are primes and $\alpha_j > 0$ for some $\ell \in \mathbb{N}$. Let $d, n \in \mathbb{N}$ such that $d \leq n$ and sample $b \leftarrow U(\mathbb{Z}_q^{\leq n+d-1}[x])$. Then,

$$\Pr[\text{rank}(\text{Hank}(b)) = d] \geq \prod_{j \in [\ell]} \left(1 - \frac{1}{p_j}\right).$$

Proof: *Case 1.* Assume that q is prime. Any Hankel matrix of order $d + n - 1$ can be represented as the matrix product of the corresponding Toeplitz matrix of order $d + n - 1$ times the non-singular reflection matrix \mathbf{J}_n of order n whose anti-diagonal elements are 1's and all other entries are 0's. Thus, the rank of a given Hankel matrix is the same as the one of the corresponding Toeplitz matrix. For the case $d = n$, it follows from [KL96, Thm. 4] that the total number of Hankel matrices of full rank d is exactly $(q - 1)q^{2d-2}$. If we choose $b \leftarrow U(\mathbb{Z}_q^{\leq n+d-1}[x])$, then

$$\Pr[\text{rank}(\text{Hank}(b)) = d] = \frac{(q - 1)q^{2d-2}}{q^{2d-1}} = 1 - \frac{1}{q}.$$

For $d < n$, the $d \times n$ Hankel matrix has full rank d if at least the left $d \times d$ submatrix, which is naturally a $d \times d$ Hankel matrix as well, has rank d . This happens with probability at least $1 - \frac{1}{q}$.

Case 2. Assume that $q = p^\alpha$ for some p prime and $\alpha > 0$. Initially, consider the case $d = n$. Any Hankel matrix \mathbf{A} can be represented as $\mathbf{A} = p\mathbf{Q} + \mathbf{R}$, where both \mathbf{R} and \mathbf{Q} are Hankel matrices with coefficients in \mathbb{Z}_p and $\mathbb{Z}_{p^{\alpha-1}}$, respectively. This formula follows from integer division by p with remainder, i.e., Euclidean division. Any element from \mathbb{Z}_{p^α} , when divided by p , has a remainder in \mathbb{Z}_p and quotient in $\mathbb{Z}_{p^{\alpha-1}}$. This representation is unique, thus preserves the structure of the matrix \mathbf{A} . Since \mathbf{A} is a Hankel matrix, so are \mathbf{Q} and \mathbf{R} . The matrix \mathbf{A} has full rank in \mathbb{Z}_{p^α} if and only if \mathbf{R} has full rank in \mathbb{Z}_p . Hence, we can deduce from the previous case that the number of Hankel matrices of full rank equals $(p - 1)p^{(\alpha-1)(2d-1)+(2d-2)}$.

If we sample $b \leftarrow U(\mathbb{Z}_q^{\leq n+d-1}[x])$, then

$$\Pr[\text{rank}(\text{Hank}(b)) = d] = \frac{(p-1)p^{(\alpha-1)(2d-1)+(2d-2)}}{p^{\alpha(2d-1)}} = 1 - \frac{1}{p}.$$

For $d < n$, using the same argument as before, the probability is at least $1 - \frac{1}{p}$.

Case 3. Let $q \in \mathbb{N}$ with unique prime power factorization given by $q = \prod_{j \in [\ell]} p_j^{\alpha_j}$, where p_j are primes and $\alpha_j > 0$ for some $\ell \in \mathbb{N}$. For the case $d = n$, it follows from the Chinese Remainder Theorem that the number of Hankel matrices of full rank d modulo q equals the product of the number of Hankel matrices of full rank d modulo $p_j^{\alpha_j}$ which is given by

$$\prod_{j \in [\ell]} (p_j - 1) p_j^{(\alpha_j - 1)(2d-1) + (2d-2)}.$$

Thus, if we sample $b \leftarrow U(\mathbb{Z}_q^{\leq n+d-1}[x])$, then $\Pr[\text{rank}(\text{Hank}(b)) = d] = \prod_{j \in [\ell]} \left(1 - \frac{1}{p_j}\right)$. Similarly as before, for $d < n$ and $b \leftarrow U(\mathbb{Z}_q^{\leq n+d-1}[x])$, then

$$\Pr[\text{rank}(\text{Hank}(b)) = d] \geq \prod_{j \in [\ell]} \left(1 - \frac{1}{p_j}\right).$$

We denote by $(\mathbb{Z}_q^{\leq n+d-1}[x])^\times$ the set of polynomials of $\mathbb{Z}_q^{\leq n+d-1}[x]$ with full rank Hankel matrix. As said earlier, Hankel matrices play an important role in the middle-product setting, as they can be used to represent the middle-product of two polynomials. More precisely, for $a \in \mathbb{Z}_q^{\leq n}[x]$ and $b \in \mathbb{Z}_q^{\leq n+d-1}[x]$, their middle-product can be written as a matrix-vector product

$$a \odot_d b = \text{Hank}(b) \cdot \text{rev}(\mathbf{a}),$$

where $\text{rev}(\mathbf{a})$ denotes the coefficient vector of a in reverse order. We write MP-LWE^\times to denote the problem MP-LWE restricted to secrets from the set $(\mathbb{Z}_q^{\leq n+d-1}[x])^\times$.

Lemma 4.2 (Uniformity of the Middle-Product)

Let $d, n \in \mathbb{N}$ such that $d \leq n$ and let b be a fixed element of $(\mathbb{Z}_q^{\leq n+d-1}[x])^\times$. If we sample $a \leftarrow U(\mathbb{Z}_q^{\leq n}[x])$, then $a \odot_d b$ is uniformly random in $\mathbb{Z}_q^{\leq d}[x]$.

Proof: We can write $a \odot_d b = \text{Hank}(b) \cdot \text{rev}(\mathbf{a})$. For any $d \leq n$ and full rank matrix $\mathbf{A} \in \mathbb{Z}_q^{d \times n}$, the mapping from \mathbb{Z}_q^n to \mathbb{Z}_q^d given by multiplication with \mathbf{A} is surjective. As a is chosen uniformly at random and the Hankel matrix of b has full rank d , their middle-product is also uniformly distributed. ■

4.3 Middle-Product Learning With Rounding (MP-LWR)

In the following, we define a new hardness assumption that we call Middle-Product Computational Learning With Rounding (MP-CLWR) and which is an adaption of the Ring Computational Learning With Rounding (R-CLWR) assumption by Chen et al. [CZZ18] to the middle-product setting. As explained before, the motivation to introduce the problem in its *computational* form stems from the fact that it possesses an efficient reduction from decision MP-LWE with a polynomially large modulus and at the same time we can use it to build a provably secure encryption scheme (see Chapter 6). For a detailed introduction and motivation of the computational notion, see [CZZ18, Section 3].

In order to define this computational assumption, we need to introduce our experiment setting, illustrated in Protocol 4.1. Within the experiment, three different parties in form of algorithms appear: A challenger \mathcal{C} interacting with an adversary \mathcal{A} who is receiving their samples from a source \mathcal{S} . All three algorithms are restricted to be probabilistic and polynomial-time (PPT). As a first step, the source \mathcal{S} generates a sample (X, aux) using two sets called var and con . They then send this sample to the challenger \mathcal{C} , who computes, with the help of this sample, a tuple $(\text{Input}, \text{Target})$. The adversary only receives the Input part of the tuple to compute the Output . The adversary wins the experiment if Output equals Target .

Protocol 4.1: The general experiment setting

$\text{Exp}(\mathcal{C}, \mathcal{A}, \mathcal{S})$
1 : $(X, \text{aux}) \leftarrow \mathcal{S}(\text{var}, \text{con})$
2 : $(\text{Input}, \text{Target}) \leftarrow \mathcal{C}(X, \text{aux})$
3 : $\text{Output} \leftarrow \mathcal{A}(\text{Input})$
4 : return $\text{Output} = \text{Target}$

The idea of the computational assumption is to consider two different experiments with the same challenger \mathcal{C} and adversary \mathcal{A} but with different sources \mathcal{S}_1 and \mathcal{S}_2 , which differ in the distribution var but have the same distribution con , motivating the notation var for variable and con for constant. The new notion guarantees that if \mathcal{A} cannot compute Target from X_1 generated by \mathcal{S}_1 , then they are not able to compute Target from X_2 generated by \mathcal{S}_2 either.

We illustrate the new notion in Protocol 4.2 below. Recall that we already illustrated it in a more informal way in Figure 4.1. Let \mathcal{C} be an arbitrary challenger. If the success probability of any adversary \mathcal{A} outputting the correct answer in $\text{Exp}_1(\mathcal{C}, \mathcal{A}, \mathcal{S}_1)$ is negligible, then it is in $\text{Exp}_2(\mathcal{C}, \mathcal{A}, \mathcal{S}_2)$ as well.

Protocol 4.2: The experiment setting of the computational assumption

$\text{Exp}_1(\mathcal{C}, \mathcal{A}, \mathcal{S}_1)$	$\text{Exp}_2(\mathcal{C}, \mathcal{A}, \mathcal{S}_2)$
1 : $(X_1, \text{aux}) \leftarrow \mathcal{S}_1(\text{var}_1, \text{con})$	1 : $(X_2, \text{aux}) \leftarrow \mathcal{S}_2(\text{var}_2, \text{con})$
2 : $(\text{Input}_1, \text{Target}_1) \leftarrow \mathcal{C}(X_1, \text{aux})$	2 : $(\text{Input}_2, \text{Target}_2) \leftarrow \mathcal{C}(X_2, \text{aux})$
3 : $\text{Output}_1 \leftarrow \mathcal{A}(\text{Input}_1)$	3 : $\text{Output}_2 \leftarrow \mathcal{A}(\text{Input}_2)$
4 : return $\text{Output}_1 = \text{Target}_1$	4 : return $\text{Output}_2 = \text{Target}_2$

Now, we define our new MP-CLWR assumption which is an adaption of the R-CLWR assump-

tion from [CZZ18] to the middle-product setting. As an analog of the notion of units in the original paper, we define $(\mathbb{Z}_q^{\leq n+d-1}[x])^\times$ as the set of all polynomials over \mathbb{Z}_q having degree less than $n+d-1$ and a Hankel matrix of order $d \times n$ of full rank d . Such unit elements are needed to map a uniform random element again to a uniform random element. Lemma 4.2 shows that the same property holds for elements with a full rank Hankel matrix.

The integers d and n define the parameters of the middle-product, q defines the general and p the rounding modulus. The number of samples has to be fixed beforehand and is given by t . For any distribution \mathcal{X} , we denote by \mathcal{X}^t the distribution that is defined by t copies of it. Thus, an instance of the latter is composed of t independent instances of \mathcal{X} .

Definition 4.3 (MP-CLWR Assumption): Let $d, n, p, q, t \in \mathbb{N}$ such that $d \leq n$ and $q \geq p \geq 2$. Sample $s \leftarrow U((\mathbb{Z}_q^{\leq n+d-1}[x])^\times)$. Denote by \mathcal{X}_s the distribution of $(a, \lfloor a \odot_d s \rfloor_p)$, where $a \leftarrow U(\mathbb{Z}_q^{\leq n}[x])$, and denote by \mathcal{U} the distribution of $(a, \lfloor b \rfloor_p)$, where $a \leftarrow U(\mathbb{Z}_q^{\leq n}[x])$ and $b \leftarrow U(\mathbb{Z}_q^{\leq d}[x])$. For $j \in \{1, 2\}$ define the input for the source \mathcal{S}_j as $(\text{var}_j, \text{con})$, where var_1 denotes the distribution \mathcal{X}_s^t , and var_2 the distribution \mathcal{U}^t , and con is an arbitrary distribution over $\{0, 1\}^*$ which is independent from var_1 and var_2 . For a fixed challenger \mathcal{C} let $\mathcal{P}_{\mathcal{C}, \mathcal{A}}$ be the probability for an adversary \mathcal{A} to win $\text{Exp}_1(\mathcal{C}, \mathcal{A}, \mathcal{S}_1)$, while $\mathcal{Q}_{\mathcal{C}, \mathcal{A}}$ be that for \mathcal{A} to win $\text{Exp}_2(\mathcal{C}, \mathcal{A}, \mathcal{S}_2)$.

The MP-CLWR $_{p,q,n,d,t}$ assumption states that for any challenger \mathcal{C} if $\mathcal{Q}_{\mathcal{C}, \mathcal{A}}$ is negligible for any adversary \mathcal{A} , so is $\mathcal{P}_{\mathcal{C}, \mathcal{A}}$. We call the difference $|\mathcal{P}_{\mathcal{C}, \mathcal{A}} - \mathcal{Q}_{\mathcal{C}, \mathcal{A}}|$ the advantage of the adversary \mathcal{A} .

Correspondingly, we also define the Middle-Product Computational Rounded Learning With Errors (MP-CRLWE) assumption which is important in the hardness reduction in Section 4.4.

Definition 4.4 (MP-CRLWE Assumption): Let $d, n, p, q, t \in \mathbb{N}$ such that $d \leq n$ and $q \geq p \geq 2$. Sample $s \leftarrow U((\mathbb{Z}_q^{\leq n+d-1}[x])^\times)$. Let ψ be an error distribution over $\mathbb{R}^{\leq d}[x]$. Denote by $\mathcal{Y}_{s,\psi}$ the distribution of $(a, \lfloor a \odot_d s + e \rfloor_p)$, where $a \leftarrow U(\mathbb{Z}_q^{\leq n}[x])$ and $e \leftarrow \psi$ and denote by \mathcal{U} the distribution of $(a, \lfloor b \rfloor_p)$, where $a \leftarrow U(\mathbb{Z}_q^{\leq n}[x])$ and $b \leftarrow U(\mathbb{Z}_q^{\leq d}[x])$. For $j \in \{1, 2\}$ define the input for \mathcal{S}_j as $(\text{var}_j, \text{con})$, where var_1 denotes the distribution $\mathcal{Y}_{s,\psi}^t$, and var_2 the distribution \mathcal{U}^t , and con is an arbitrary distribution over $\{0, 1\}^*$ which is independent from var_1 and var_2 . For a fixed challenger \mathcal{C} let $\mathcal{P}'_{\mathcal{C}, \mathcal{A}}(\psi)$ be the probability for an adversary \mathcal{A} to win $\text{Exp}_1(\mathcal{C}, \mathcal{A}, \mathcal{S}_1)$, while $\mathcal{Q}_{\mathcal{C}, \mathcal{A}}$ be that for \mathcal{A} to win $\text{Exp}_2(\mathcal{C}, \mathcal{A}, \mathcal{S}_2)$.

The MP-CRLWE $_{p,q,n,d,t,\psi}$ assumption related to the error distribution ψ states that for any challenger \mathcal{C} if $\mathcal{Q}_{\mathcal{C}, \mathcal{A}}$ is negligible for any adversary \mathcal{A} , so is $\mathcal{P}'_{\mathcal{C}, \mathcal{A}}(\psi)$. We call the difference $|\mathcal{P}'_{\mathcal{C}, \mathcal{A}}(\psi) - \mathcal{Q}_{\mathcal{C}, \mathcal{A}}|$ the advantage of the adversary \mathcal{A} .

4.4 Hardness of Computational MP-LWR

We now prove the hardness of MP-CLWR with the help of a reduction from the decision MP-LWE problem to the MP-CLWR problem. The decision version of MP-LWE itself can be reduced from decision P-LWE for an exponentially large class of defining polynomials [RSSS17]. As P-LWE benefits from worst-case to average-case reductions from lattice problems, our new MP-CLWR assumption also enjoys worst-case hardness.

In the following we need the notion of B -bounded and balanced distributions. More precisely,

a distribution ψ is B -bounded with probability at least $\delta \in [0, 1]$ for a real number $B \geq 0$, if $\Pr_{x \leftarrow \psi}[|x| \leq B] \geq \delta$. We further call a B -bounded distribution ψ balanced if $\Pr_{x \leftarrow \psi}[|x| \leq 0] \geq 1/2$ and at the same time $\Pr_{x \leftarrow \psi}[|x| \geq 0] \geq 1/2$.

Theorem 4.1 (Hardness of MP-CLWR)

Let $d, n, p, q, t \in \mathbb{N}$ such that $d \leq n$ and $q \geq p \geq 2$. Further, let $q = \prod_{j \in [\ell]} p_j^{\alpha_j}$ be the prime power factorization of q with some $\ell \in \mathbb{N}$, where p_j is prime and $\alpha_j > 0$ for all $j \in [\ell]$. Let ψ 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 decision $\text{MP-LWE}_{q,n,d,\psi}$ to $\text{MP-CLWR}_{p,q,n,d,t}$, with t the number of samples fixed beforehand.

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

$$\left(\frac{1}{e^2} (\varepsilon + \mathcal{Q}_{\mathcal{C},\mathcal{A}})^2 \right) \cdot \prod_{j \in [\ell]} \left(1 - \frac{1}{p_j} \right),$$

where e denotes the Euler's number and $\mathcal{Q}_{\mathcal{C},\mathcal{A}}$ is as in Definition 4.3.

Our reduction from MP-LWE to MP-CLWR preserves the dimension and the modulus, working for any polynomial-sized modulus q . To prove Theorem 4.1, we need the following sequence of reductions, as illustrated in Figure 4.2.




Figure 4.2: Overview of the proof of Theorem 4.1.

The first part of this sequence, Lemma 4.3, is a standard reduction from decision MP-LWE to decision MP-LWE^{\times} , where the latter one denotes the MP-LWE problem with secrets sampled uniformly at random from the set of elements having a full rank Hankel matrix. Lemma 4.4 maps instances of MP-LWE^{\times} to instances of MP-CRLWE by simply rounding the second part of any instance. Our results on random Hankel matrices from Section 4.2 are used in Lemma 4.5, where we show a reduction from the rounded middle-product LWE problem to the middle-product LWR problem, for their computational versions. Note that this reduction uses the Rényi divergence and thus asks for fixing the requested number of samples t a priori. This is a necessary requirement which is also imposed in previous works [BGM⁺16, CZZ18].

Lemma 4.3 (MP-LWE to MP-LWE[×])

Let $d, n, q \in \mathbb{N}$ such that $d \leq n$. Let ψ be an error distribution over $\mathbb{R}^{<d}[x]$. Further, let $q = \prod_{j \in [\ell]} p_j^{\alpha_j}$ be the prime power factorization of q with some $\ell \in \mathbb{N}$, where p_j is prime and $\alpha_j > 0$ for all $j \in [\ell]$. If there is a PPT algorithm solving MP-LWE[×] $_{q,n,d,\psi}$ with non-negligible advantage ε , then there is a PPT algorithm solving MP-LWE $_{q,n,d,\psi}$ with non-negligible advantage at least

$$\varepsilon \cdot \prod_{j \in [\ell]} \left(1 - \frac{1}{p_j}\right).$$

Proof: Let $(a_k, b_k)_k$ be arbitrarily many input samples of MP-LWE $_{q,n,d,\psi}$, where the instances either come from the uniform distribution over $\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x]$ or from the MP-LWE distribution, i.e., $b_k = a_k \odot_d s + e_k$, where $s \leftarrow U(\mathbb{Z}_q^{n+d-1}[x])$ and $e_k \leftarrow \psi$ for all k . A solver can take the samples $(a_j, b_j)_j$ and query an oracle of MP-LWE[×] $_{q,n,d,\psi}$ on it. If the instances come from the uniform distribution, the oracle succeeds as usual. If the instances come from the real MP-LWE distribution, the oracle surely succeeds only if the Hankel matrix of s has full rank d . As shown in Lemma 4.1, the probability that the Hankel matrix of s has full rank d is at least $\prod_{j \in [\ell]} \left(1 - \frac{1}{p_j}\right)$. Assuming that the oracle succeeds with non-negligible probability ε in general, it now succeeds with probability at least $\varepsilon \cdot \prod_{j \in [\ell]} \left(1 - \frac{1}{p_j}\right)$, which completes the proof. ■

The second lemma is an adaption of [CZZ18, Lem. 12] to our context.

Lemma 4.4 (MP-LWE[×] to MP-CRLWE)

Let $d, n, q, p, t \in \mathbb{N}$ such that $d \leq n$ and $q \geq p \geq 2$. Let ψ be an error distribution over $\mathbb{R}^{<d}[x]$. Assume that the advantage of any PPT algorithm to solve the decision MP-LWE[×] $_{q,n,d,\psi}$ problem is less than ε , then we have

$$|\mathcal{P}'_{\mathcal{C},\mathcal{A}}(\psi) - \mathcal{Q}_{\mathcal{C},\mathcal{A}}| < \varepsilon,$$

for any PPT adversary \mathcal{A} and PPT challenger \mathcal{C} . Thus, there is a reduction from MP-LWE[×] $_{q,n,d,\psi}$ to MP-CRLWE $_{p,q,n,d,t,\psi}$, with t the number of samples fixed beforehand.

Proof: In order to show this reduction, we construct an adversary \mathcal{B} to solve the decision MP-LWE[×] $_{q,n,d,\psi}$ problem. This adversary \mathcal{B} plays at the same time the role of the challenger \mathcal{C} in the MP-CRLWE experiment. At the beginning, \mathcal{B} receives a tuple of samples $(x_j, y_j)_{j \in [t]}$. They set $a_j = x_j$ and $b_j = \lfloor y_j \rfloor_p$ for all $j \in [t]$ and $X = (a_j, b_j)_{j \in [t]}$. As a challenger of the experiment, \mathcal{B} can compute the corresponding Input and Target. They send the input to some adversary \mathcal{A} and verify if the Output of \mathcal{A} equals the Target. If this is the case, \mathcal{B} outputs 1, otherwise 0.

If $(x_j, y_j)_{j \in [t]}$ are MP-LWE samples, then are $(a_j, b_j)_{j \in [t]}$ samples from $\mathcal{Y}_{s,\psi}$, used in the MP-CRLWE assumption. Thus, $\Pr[\mathcal{B}((x_j, y_j)_j) = 1] = \mathcal{P}'_{\mathcal{C},\mathcal{A}}(\psi)$. On the other hand, if $(x_j, y_j)_j$

are uniform samples, then are $(a_j, b_j)_j$ also uniformly distributed. Hence, $\Pr[\mathcal{B}((x_j, y_j)_j) = 1] = \mathcal{Q}_{\mathcal{C}, \mathcal{A}}$.

Note that the **(Input, Target)** pair can be seen as a function in the input X . The maximal information on X it can contain, is if \mathcal{B} sets **Input** = X (and as a target the require to distinguish the used distribution). Assuming the hardness of decision MP-LWE, we have $|\mathcal{P}'_{\mathcal{C}, \mathcal{A}}(\psi) - \mathcal{Q}_{\mathcal{C}, \mathcal{A}}| < \varepsilon$, for negligible ε and for any PPT adversary \mathcal{A} . In particular, the MP-CRLWE assumption holds: If $\mathcal{Q}_{\mathcal{C}, \mathcal{A}}$ is negligible, so is $\mathcal{P}'_{\mathcal{C}, \mathcal{A}}$ for the same challenger \mathcal{C} and adversary \mathcal{A} using the equation above. ■

The third and last reduction is an adaption of Lemma 8 and Lemma 9 in [CZZ18], based on the results of [BGM⁺16], together with our results about random Hankel matrices of Section 4.2.

Lemma 4.5 (MP-CRLWE to MP-CLWR)

Let $d, n, q, p, t \in \mathbb{N}$ such that $d \leq n$ and $q \geq p \geq 2$. Further, sample $s \leftarrow U((\mathbb{Z}_q^{\leq n+d-1}[x])^\times)$, and let \mathcal{X}_s and \mathcal{Y}_s denote the random variables of a single MP-CLWR sample $(a, \lfloor a \odot_d s \rfloor_p)$ and a single MP-CRLWE sample $(a, \lfloor a \odot_d s + e \rfloor_p)$, respectively. Further, let ψ be an error distribution over $\mathbb{R}^{\leq d}[x]$ which is balanced and B -bounded with probability at least δ , where $q > 2pBdt$ and $\delta \geq 1 - \frac{1}{td}$. Then we have

$$(\mathcal{P}_{\mathcal{C}, \mathcal{A}})^2 \leq \mathcal{P}'_{\mathcal{C}, \mathcal{A}}(\psi) \cdot e^2,$$

where e is the Euler's number. Hence, there is a reduction from $\text{MP-CRLWE}_{p,q,n,d,t,\psi}$ to $\text{MP-CLWR}_{p,q,n,d,t}$.

Proof: Let $s \leftarrow U((\mathbb{Z}_q^{\leq n+d-1}[x])^\times)$ be fixed throughout the proof. Using the multiplicativity and probability preservation properties of the Rényi divergence, see Lemma 1.20, we have $(\mathcal{P}_{\mathcal{C}, \mathcal{A}})^2 \leq \mathcal{P}'_{\mathcal{C}, \mathcal{A}}(\psi) \cdot \text{RD}_2(\mathcal{X}_s \| \mathcal{Y}_s)^t$. In the following we show that the Rényi divergence of \mathcal{X}_s and \mathcal{Y}_s can be bounded above by

$$\text{RD}_2(\mathcal{X}_s \| \mathcal{Y}_s) \leq \frac{(1 + 2pB/q)^d}{\delta^d}.$$

Following the definition of the Rényi divergence it yields

$$\begin{aligned} \text{RD}_2(\mathcal{X}_s \| \mathcal{Y}_s) &= E_{a \leftarrow U(\mathbb{Z}_q^{\leq n}[x])} \frac{\Pr \left[\mathcal{X}_s = (a, \lfloor a \odot_d s \rfloor_p) \right]}{\Pr \left[\mathcal{Y}_s = (a, \lfloor a \odot_d s \rfloor_p) \right]} \\ &= E_{a \leftarrow U(\mathbb{Z}_q^{\leq n}[x])} \frac{1}{\Pr_{e \leftarrow \psi} \left[\lfloor a \odot_d s + e \rfloor_p = \lfloor a \odot_d s \rfloor_p \right]}, \end{aligned}$$

where we use that the expected value over $x \leftarrow \mathcal{X}_s$ is equivalently defined by the expected value over $a \leftarrow U(\mathbb{Z}_q^{\leq n}[x])$. First, we define the border elements in \mathbb{Z}_q with respect to B and p

by

$$\text{Bor}_{p,q}(B) = \left\{ x \in \mathbb{Z}_q : \lfloor x + B \rfloor_p \neq \lfloor x \rfloor_p \right\}.$$

These are the elements in \mathbb{Z}_q which are close to the rounding boundary. It yields $|\text{Bor}_{p,q}(B)| \leq 2Bp$. For $0 \leq t \leq d$ we define

$$\text{Bad}_{s,t} = \left\{ a \in \mathbb{Z}_q^{\leq n}[x] : |\{j \in [d] : (a \odot_d s)_j \in \text{Bor}_{p,q}(B)\}| = t \right\},$$

where $(a \odot_d s)_j$ denotes the j -th coefficient of $a \odot_d s$.^a In other words, $\text{Bad}_{s,t}$ defines for a given polynomial s and number of coefficients t , the set of polynomials a in $\mathbb{Z}_q^{\leq n}[x]$ such that the middle-product $a \odot_d s$ has exactly t coefficients close to the rounding boundary. Now we fix t and assume $a \in \text{Bad}_{s,t}$. For any $j \in [d]$ with $(a \odot_d s)_j \notin \text{Bor}_{p,q}(B)$, it yields

$$\Pr_{(e_k)_{k=e \leftarrow \psi}} \left[\lfloor (a \odot_d s)_j + e_j \rfloor_p = \lfloor (a \odot_d s)_j \rfloor_p \right] \geq \delta,$$

as e_j is the j -th coefficient of e , where e is sampled from the distribution ψ which is B -bounded with probability at least δ . If $(a \odot_d s)_j \in \text{Bor}_{p,q}(B)$, then

$$\Pr_{(e_k)_{k=e \leftarrow \psi}} \left[\lfloor (a \odot_d s)_j + e_j \rfloor_p \neq \lfloor (a \odot_d s)_j \rfloor_p \right] \geq \frac{1}{2},$$

because e_j is the j -th coefficient of e , where e is sampled from a balanced distribution. Thus, the probabilities of $e_j \in [-B, 0]$ or in $[0, B]$ are each greater or equal to $\frac{1}{2}$ and $\lfloor (a \odot_d s)_j + e_j \rfloor_p \neq \lfloor (a \odot_d s)_j \rfloor_p$ happens in exactly one of the two cases. Since each coefficient of e is independently distributed and $a \odot_d s$ has exactly t coefficients in $\text{Bor}_{p,q}(B)$, it yields

$$\Pr_{e \leftarrow \psi} \left[\lfloor a \odot_d s + e \rfloor_p = \lfloor a \odot_d s \rfloor_p \right] \geq \frac{1}{2^t} \cdot \delta^{d-t} \geq \frac{1}{2^t} \cdot \delta^d.$$

By Lemma 4.2, we know that if a is uniform in $\mathbb{Z}_q^{\leq n}[x]$, so is $a \odot_d s \in \mathbb{Z}_q^{\leq d}[x]$ (for $s \in (\mathbb{Z}_q^{\leq n+d-1}[x])^\times$). Thus, it yields

$$\Pr [a \in \text{Bad}_{s,t}] \leq \binom{d}{t} \left(1 - \frac{|\text{Bor}_{p,q}(B)|}{q} \right)^{d-t} \left(\frac{|\text{Bor}_{p,q}(B)|}{q} \right)^t.$$

Hence,

$$\begin{aligned} \text{RD}_2(\mathcal{X}_s \| \mathcal{Y}_s) &\leq \delta^{-d} \sum_{t \in [d]} 2^t \cdot \Pr [a \in \text{Bad}_{s,t}] \\ &= \delta^{-d} \sum_{t \in [d]} \binom{d}{t} \left(1 - \frac{|\text{Bor}_{p,q}(B)|}{q} \right)^{d-t} \left(2 \cdot \frac{|\text{Bor}_{p,q}(B)|}{q} \right)^t \\ &= \delta^{-d} \left(1 + \frac{|\text{Bor}_{p,q}(B)|}{q} \right)^d \\ &\leq \delta^{-d} \left(1 + \frac{2pB}{q} \right)^d, \end{aligned}$$

where we used in the third equation the standard binomial theorem. From the results above, we can derive

$$\text{RD}_2(\mathcal{X}_s \| \mathcal{Y}_s)^t \leq \frac{(1 + 2pB/q)^{td}}{\delta^{td}} \leq \frac{(1 + 1/td)^{td}}{(1 - 1/td)^{td}} \approx e^2,$$

where $\delta \geq 1 - \frac{1}{td}$ and $q > 2pBdt$. ■

^aBy abuse of notation we write $j \in [d]$, even though strictly speaking the coefficients run from 0 to $d - 1$, and not from 1 to d .

Putting Lemma 4.3, Lemma 4.4 and Lemma 4.5 together completes the proof of Theorem 4.1.

Chapter 5

Partial Vandermonde Problems

The content of this chapter is based on a joint work with Amin Sakzad and Ron Steinfeld, which was initiated during a research stay at the Monash University in Melbourne from October to December 2019. The results haven't been published before and are publicly presented for the first time in this thesis.

Contents

5.1	Introduction	89
5.1.1	Our Contributions	90
5.2	Definition of Partial Vandermonde Problems	91
5.2.1	Partial Vandermonde Knapsack Problem (PV-Knap)	91
5.2.2	Partial Vandermonde Learning With Errors (PV-LWE)	93
5.2.3	PASS Problem	94
5.3	Equivalence of PV-Knap and PV-LWE	95
5.4	Efficient Computation of Partial Vandermonde Transform	97
5.4.1	Computation of the Partial Vandermonde Transformation	98
5.4.2	Computation of the Transposed Partial Vandermonde Transformation	98

5.1 Introduction

As elaborated in the introduction of this manuscript, lattice-based cryptography is a relatively young research field of public key cryptography that was initiated at the end of the 1990s by two different branches. On the one hand, there have been proposals benefiting from strong theoretical connections to presumed hard worst-case lattice problems [Ajt96, AD97], leading to the development of public key cryptography based on the SIS and LWE problems, see for instance the survey by Peikert [Pei16a]. Both, the module variant of LWE that we study in Chapter 2 and Chapter 3, and the middle-product variant of LWR from Chapter 4, can be seen as descendants of Ajtai's work, still benefiting from worst-case to average-case reductions to well-studied problems on (structured) lattices. On the other hand, however, very efficient schemes basing their security on average-case structured lattice problems have been introduced, the most popular among them is the NTRU encryption scheme by Hoffstein et al. [HPS98]. In this chapter, we study a problem which can be seen as the offspring of the second line of works.

Following the latter approach, Hoffstein et al. [HPS⁺14] propose a digital signature scheme called PASS Sign, whose security is based on the difficulty of recovering a polynomial of small

norm having access only to a partial list of its discrete Fourier transform. Later, Lu et al. [LZA18] complement the proposal by moving from the partial Fourier transform (evaluation at all roots of unity) to the partial Vandermonde transform (evaluation only at the *primitive* roots of unity) and by giving a rigorous proof of security.

The problem that underlies PASS Sign, as given in [LZA18], can be presented as follows. Let q be a prime, $R = \mathbb{Z}[x]/\langle f(x) \rangle$ denote the ring of integers of the ν -th cyclotomic number field of degree $\deg(f(x)) = n$, such that $f(x)$ splits into linear factors mod q . More precisely, $f(x) = \prod_{j \in [n]} (x - \omega_j) \pmod q$, where $\{\omega_j\}_{j \in [n]}$ are the ν -th primitive roots of unity in \mathbb{Z}_q . Let $\bar{\mathbf{V}} = (\omega_j^{k-1})_{j,k \in [n]} \in \mathbb{Z}_q^{n \times n}$ denote the discrete Vandermonde matrix for $\{\omega_j\}_{j \in [n]}$, as introduced in Section 1.1.4. We denote by $\bar{\mathbf{V}}_\Omega \in \mathbb{Z}_q^{t \times n}$ the partial Vandermonde matrix consisting of $t \leq n$ subrows of $\bar{\mathbf{V}}$ specified by a random subset of t roots $\Omega \subseteq \{\omega_j\}_{j \in [n]}$. Let \mathbf{f} be a ring element of small norm, sampled from some distribution ψ . Given $\bar{\mathbf{V}}_\Omega \cdot \mathbf{f} \pmod q$, the problem asks to find \mathbf{f} . We call this the Partial Vandermonde Knapsack problem (PV-Knap).¹ A related problem is Partial Vandermonde SIS (PV-SIS), where given $\bar{\mathbf{V}}_\Omega$ one asks to find a ring element \mathbf{f} of small norm such that $\bar{\mathbf{V}}_\Omega \cdot \mathbf{f} = \mathbf{0} \pmod q$. This can be formulated as a Shortest Vector Problem (SVP) over a structured lattice.

Shortly afterwards, Hoffstein and Silverman [HS15] introduce PASS Encrypt, a Public Key Encryption (PKE) scheme whose computational building blocks are closely related to the ones of PASS Sign. It is very efficient and fulfills additive and multiplicative homomorphic properties. The algebraic structure and homomorphic properties of PASS Encrypt and the underlying partial Vandermonde problems, make them a natural starting point for the design of efficient cryptographic primitives. For example, such properties are recently exploited in the context of PASS Sign to construct compact aggregate signature schemes [DHSS20], and it is plausible that combining PASS Encrypt with PASS Sign may form the basis for various compact and efficient privacy-preserving primitives such as group signatures. Unfortunately, a main problem with PASS Encrypt to date is that its security is not well understood, no proof of security was given in [HS15] with respect to the hardness of explicit computational problems, and the scheme is deterministic and hence does not satisfy the standard notion of IND-CPA security.

5.1.1 Our Contributions

In this chapter, we make progress towards understanding the hardness assumptions needed to prove the security of PASS Encrypt. We first present the Partial Vandermonde Knapsack problem (PV-Knap) in its search and decision variant, as studied by Lu et al. [LZA18] in the context of PASS Sign.² It makes use of the partial discrete Vandermonde matrix over \mathbb{Z}_q , where q is a prime. We emphasize its connection to (average-case) ideal lattices, even though we can't show a worst-case to average-case reduction as for structured variants of LWE. We enlarge the landscape of problems that use the partial Vandermonde matrix by defining a new variant of LWE, called Partial Vandermonde Learning With Errors (PV-LWE). Later, we show the equivalence of PV-Knap and PV-LWE by exploiting the same duality connection as we have for standard Knapsack problems and LWE. As explained in the introduction of this chapter, our main motivation is to provide a security proof for PASS Encrypt. To this end, we define a variant of PV-Knap, that we call the PASS problem. This problem serves (together with the decision version of PV-Knap) as

¹We remark that $\bar{\mathbf{V}}_\Omega \cdot \mathbf{f} \pmod q$ is the vector of evaluations $(\mathbf{f}(\omega_j))_{\omega_j \in \Omega} \pmod q$ of the polynomial \mathbf{f} at the roots in Ω ; the full vector of evaluations $(\mathbf{f}(\omega_j))_{j \in [n]}$ is also known as the Number Theoretic Transform (NTT) of \mathbf{f} and $\mathbf{f}(\omega_j)$ is sometimes referred to as the j -th NTT *slot* or NTT *coefficient* of \mathbf{f} .

²Note that the name of the problem varies in the former literature between Partial Fourier Recovery problem, Vandermonde (I)SIS and Fourier (I)SIS. We think, that it's most related to the Knapsack problem and thus call it that way.

the underlying hardness assumption for (a slightly modified version of) PASS Encrypt. We present the scheme together with the security proof in Chapter 7. We conclude the chapter by showing that for a special choice of the partial Vandermonde matrix, one can accelerate computations when working with PV-Knap and PV-LWE.

5.2 Definition of Partial Vandermonde Problems

We now define the problems Partial Vandermonde Knapsack, Partial Vandermonde SIS, Partial Vandermonde LWE and a leaky variant of Partial Vandermonde Knapsack, called PASS.

5.2.1 Partial Vandermonde Knapsack Problem (PV-Knap)

The Partial Vandermonde Knapsack problem (PV-Knap) was first introduced by Hoffstein et al. [HPS⁺14].³ We start this section by defining the notion of a partial Vandermonde matrix, where we consider the discrete Vandermonde matrix $\bar{\mathbf{V}}$ as presented in Section 1.1.4. Informally speaking, the partial Vandermonde matrix is composed of a subset of rows of the full matrix $\bar{\mathbf{V}}$.

We now state the formal definition. For $\nu \in \mathbb{N}$, let $K = \mathbb{Q}[x]/\langle f(x) \rangle$ be the ν -th cyclotomic number field of degree $\deg(f(x)) = \varphi(\nu) = n$ and let $R = \mathbb{Z}[x]/\langle f(x) \rangle$ be its ring of integers. Further, let $q \in \mathbb{N}$ be a prime such that $q \equiv 1 \pmod{\nu}$. In this case, $f(x)$ completely splits in $\mathbb{Z}_q[x]$, i.e., $f(x) = \prod_{j \in [n]} (x - \omega_j)$, where every ω_j is a distinctive primitive ν -th root of unity in \mathbb{Z}_q . Furthermore, the ideal generated by q in R can be written as $\langle q \rangle = \prod_{j \in [n]} \langle q, x - \omega_j \rangle$. To ease readability, we represent in this chapter a ring element $a \in R$ directly by its coefficient vector $\mathbf{a} \in \mathbb{Z}^n$ and implicitly assume that this coefficient vector is obtained by applying the coefficient embedding τ as introduced in Section 1.1.3.

Definition 5.1 (The partial Vandermonde transform): For ν, n, q as above, let $\{\omega_j\}_{j \in [n]}$ be the set of primitive ν -th roots of unity in \mathbb{Z}_q . We divide the set $\{\omega_j\}_{j \in [n]}$ into two disjoint subsets Ω and Ω^c of size $|\Omega| = t$ and $|\Omega^c| = n - t$. The partial Vandermonde transformation matrix $\bar{\mathbf{V}}_\Omega \in \mathbb{Z}_q^{t \times n}$ and its complement $\bar{\mathbf{V}}_{\Omega^c} \in \mathbb{Z}_q^{(n-t) \times n}$ are given by

$$\bar{\mathbf{V}}_\Omega = \begin{bmatrix} 1 & \omega_{i_1} & \cdots & \omega_{i_1}^{n-1} \\ \vdots & \vdots & & \vdots \\ 1 & \omega_{i_t} & \cdots & \omega_{i_t}^{n-1} \end{bmatrix}, \quad \text{and} \quad \bar{\mathbf{V}}_{\Omega^c} = \begin{bmatrix} 1 & \omega_{i_{t+1}} & \cdots & \omega_{i_{t+1}}^{n-1} \\ \vdots & \vdots & & \vdots \\ 1 & \omega_{i_n} & \cdots & \omega_{i_n}^{n-1} \end{bmatrix},$$

where $\omega_j \in \Omega$ for $j \in [t]$ and $\omega_{i_{t+k}} \in \Omega^c$ for $k \in [n - t]$.

By rearranging and merging the rows of the two matrices $\bar{\mathbf{V}}_\Omega$ and $\bar{\mathbf{V}}_{\Omega^c}$, we obtain the full discrete Vandermonde matrix $\bar{\mathbf{V}}$. In order to get more familiar with the concept of the partial Vandermonde transform, we now provide a concrete example in small dimension.

Example 5.1 (Partial Vandermonde Matrix in Small Dimension)

Let $\nu = 16$, defining the 16-th cyclotomic field $K = \mathbb{Q}[x]/\langle x^n + 1 \rangle$ of degree $n = 8$. Further, let $q = 17$ be a prime fulfilling the condition $q \equiv 1 \pmod{\nu}$. The set of all primitive ν -th roots of unity is given by $\{3, 5, 6, 7, 10, 11, 12, 14\}$. We can set $\Omega = \{3, 5, 14, 12\}$ and $\Omega^c = \{10, 11, 7, 6\}$,

³Even though they originally called it the Partial Fourier Recovery problem.

defining

$$\overline{\mathbf{V}}_{\Omega} = \begin{bmatrix} 1 & 3 & 9 & 10 & 13 & 5 & 15 & 11 \\ 1 & 5 & 8 & 6 & 13 & 14 & 2 & 10 \\ 1 & 14 & 9 & 7 & 13 & 12 & 15 & 6 \\ 1 & 12 & 8 & 11 & 13 & 3 & 2 & 7 \end{bmatrix}, \quad \overline{\mathbf{V}}_{\Omega^c} = \begin{bmatrix} 1 & 10 & 15 & 14 & 4 & 6 & 9 & 5 \\ 1 & 11 & 2 & 5 & 4 & 10 & 8 & 3 \\ 1 & 7 & 15 & 3 & 4 & 11 & 9 & 12 \\ 1 & 6 & 2 & 12 & 4 & 7 & 8 & 14 \end{bmatrix}.$$

As all primitive roots are distinct, the matrices $\overline{\mathbf{V}}_{\Omega}$ and $\overline{\mathbf{V}}_{\Omega^c}$ have maximal row rank t and $n-t$, respectively. Multiplying the coefficient vector of an element $\mathbf{a} \in R$ by $\overline{\mathbf{V}}_{\Omega}$ (resp. $\overline{\mathbf{V}}_{\Omega^c}$) corresponds to the evaluation of \mathbf{a} at the points ω_j for $j \in [t]$ (resp. at the points ω_{i+k} for $k \in [n-t]$). To ease notations, we omit the product syntax \cdot and simply write $\overline{\mathbf{V}}_{\Omega}\mathbf{a}$ for the matrix-vector product.

Furthermore, $\overline{\mathbf{V}}_{\Omega}\mathbf{a}$ is a partial discrete canonical embedding vector of \mathbf{a} , as introduced in Section 1.1.4. Knowing $\overline{\mathbf{V}}_{\Omega}\mathbf{a}$ and $\overline{\mathbf{V}}_{\Omega^c}\mathbf{a}$ gives the complete discrete canonical embedding vector $\overline{\sigma}(\mathbf{a}) = \overline{\mathbf{V}}\mathbf{a} \in \mathbb{Z}_q^n$ and thus uniquely identifies it modulo q . The matrix $\overline{\mathbf{V}}_{\Omega}$ defines a ring homomorphism from R to \mathbb{Z}_q^t , where the latter is equipped with component-wise addition and multiplication, denoted by $+$ and \circ . Thus, $\overline{\mathbf{V}}_{\Omega}(\mathbf{a}_1 + \mathbf{a}_2) = (\overline{\mathbf{V}}_{\Omega}\mathbf{a}_1) + (\overline{\mathbf{V}}_{\Omega}\mathbf{a}_2)$ and $\overline{\mathbf{V}}_{\Omega}(\mathbf{a}_1 \cdot \mathbf{a}_2) = (\overline{\mathbf{V}}_{\Omega}\mathbf{a}_1) \circ (\overline{\mathbf{V}}_{\Omega}\mathbf{a}_2)$. In Section 5.4, we show how $\overline{\mathbf{V}}_{\Omega}\mathbf{a}$ can be efficiently computed using the Number Theoretic Transform (NTT) in dimension t for the case of power-of-2 cyclotomics and special choices of Ω .

We now present the Partial Vandermonde Knapsack problem in its search and its decision version. We define the set $\mathcal{P}_t = \{\Omega \subseteq \{\omega_j\}_{j \in [n]} : |\Omega| = t\}$ of all subsets of primitive ν -th roots of unity in \mathbb{Z}_q of size t . The corresponding set Ω for an instance of Partial Vandermonde Knapsack is chosen uniformly at random over \mathcal{P}_t .

Definition 5.2 (Partial Vandermonde Knapsack): Let ψ be a distribution over \mathbb{Z}^n . The Search Partial Vandermonde Knapsack problem PV-Knap $_{\psi}$ is the following: let $\mathbf{a} \leftarrow \psi$ and $\Omega \leftarrow U(\mathcal{P}_t)$; given $\mathbf{b} = \overline{\mathbf{V}}_{\Omega}\mathbf{a} \bmod q$, find the solution \mathbf{a} . The Decision Partial Vandermonde Knapsack problem dec-PV-Knap $_{\psi}$ asks to distinguish for a given tuple $(\overline{\mathbf{V}}_{\Omega}, \mathbf{b}) \in \mathbb{Z}_q^{t \times n} \times \mathbb{Z}_q^t$, where $\Omega \leftarrow U(\mathcal{P}_t)$, if the vector \mathbf{b} was sampled from the uniform distribution over \mathbb{Z}_q^t , or if the tuple is given as an PV-Knap $_{\psi}$ instance.

The problem is assumed to be hard to solve if the distribution ψ provides elements of \mathbb{Z}^n with small norms, where smallness is with respect to q . The choice of ψ regarding the other parameters also defines whether the solution to this problem is unique or not. We further define a homogeneous variant of PV-Knap with respect to the Euclidean norm.

Definition 5.3 (Partial Vandermonde SIS): Let $\beta \in \mathbb{N}$. The Partial Vandermonde SIS problem PV-SIS $_{\beta}$ consists in finding $\mathbf{a} \in R$ for a given $\Omega \leftarrow U(\mathcal{P}_t)$ such that $\overline{\mathbf{V}}_{\Omega}\mathbf{a} = \mathbf{0} \bmod q$ and $\|\mathbf{a}\|_2 \leq \beta$.

Similar to previous works [HPS⁺14, LZA18, DHSS20], we call this problem Partial Vandermonde SIS in analogy to the standard Short Integer Solution (SIS) problem, as introduced by Ajtai [Ajt96]. Here, the fully random matrix over \mathbb{Z}_q is replaced by a random structured one, the partial Vandermonde transform $\overline{\mathbf{V}}_{\Omega}$. Note that an instance of PV-SIS corresponds to an instance

of Id-SVP (Def. 1.1) in the special class of ideal lattices

$$\mathcal{I}_{\Omega,q} := \Lambda_q^\perp(\overline{\mathbf{V}}_\Omega) = \{\mathbf{a} \in R : \overline{\mathbf{V}}_\Omega \mathbf{a} = \mathbf{0} \pmod{q}\},$$

as presented in Section 1.2.1. Using the definition of $\overline{\mathbf{V}}_\Omega$, the ideal lattice $\mathcal{I}_{\Omega,q}$ can also be presented as the product of half of the factors of the ideal $\langle q \rangle$, i.e., $\mathcal{I}_{\Omega,q} = \prod_{\omega_j \in \Omega} \langle q, x - \omega_j \rangle$. The determinant of the ideal lattice equals the norm of the ideal $\mathcal{I}_{\Omega,q}$, given by q^t . Thus, for a given ring element \mathbf{a} of small norm, the partial Vandermonde transform $\overline{\mathbf{V}}_\Omega \mathbf{a} \pmod{q}$ is a way to specify the coset $\mathbf{a} + \mathcal{I}_{\Omega,q}$. Further, the complement partial Vandermonde transform $\overline{\mathbf{V}}_{\Omega^c} \mathbf{a} \pmod{q}$ specifies the coset $\mathbf{a} + \mathcal{I}'_{\Omega,q}$, where $\mathcal{I}'_{\Omega,q} := \langle q \rangle \mathcal{I}_{\Omega,q}^{-1}$. Given $\mathbf{a} + \mathcal{I}_{\Omega,q}$ and $\mathbf{a} + \mathcal{I}'_{\Omega,q}$ uniquely defines $\mathbf{a} + \langle q \rangle$.

We remark that this class of Id-SVP instances is considered for attacks by Pan et al. [PXWC21]. However, we are in the fully splitting case ($q = 1 \pmod{\nu}$) and thus their algorithm essentially needs an oracle of SVP of dimension 2^n , which makes the attack vacuous.

In the following, we show that knowing $\mathbf{b} = \overline{\mathbf{V}}_\Omega \mathbf{a} \pmod{q}$ (but not explicitly \mathbf{a}) suffices to compute $\mathbf{B} = \overline{\mathbf{V}}_\Omega \text{Rot}(\mathbf{a}) \pmod{q}$, where $\text{Rot}(\mathbf{a})$ is the matrix of multiplication as presented in Section 1.1.3. In other words, knowing \mathbf{B} does not reveal more information than knowing \mathbf{b} . The other direction is trivial, as the first column of \mathbf{B} equals \mathbf{b} .

Lemma 5.1

Given n, t, q, Ω as above, defining $\overline{\mathbf{V}}_\Omega = (\omega_{ij}^{k-1})_{j \in [t], k \in [n]}$. Let $\mathbf{b} = \overline{\mathbf{V}}_\Omega \mathbf{a} \pmod{q} \in \mathbb{Z}_q^t$ for some ring element $\mathbf{a} \in R$. Then, for $k \in [n]$, the k -th column of the matrix $\mathbf{B} = \overline{\mathbf{V}}_\Omega \text{Rot}(\mathbf{a})$ is given by $\mathbf{b} \circ (\omega_{i_1}^{k-1}, \dots, \omega_{i_t}^{k-1})^T$.

Proof: Let $a(x)$ denote the polynomial that corresponds to $\mathbf{a} \in R = \mathbb{Z}[x]/\langle f(x) \rangle$, i.e., $\tau(\mathbf{a}) = \mathbf{a}$. For $k \in [n]$, the k -th column of $\text{Rot}(\mathbf{a})$ is given by the coefficient vector of $a(x) \cdot x^{k-1} \pmod{f(x)}$. For $\ell \in [t]$, the ℓ -th coefficient of \mathbf{b} corresponds to the evaluation of $a(x)$ at $\omega_{i_\ell} \in \Omega$. We know that $\mathbf{B} = (b_{\ell k})_{\ell \in [t], k \in [n]}$ is given by $b_{\ell k} = (a(x) \cdot x^{k-1})(\omega_{i_\ell}) = a(\omega_{i_\ell}) \cdot \omega_{i_\ell}^{k-1}$. The homomorphic properties of $\overline{\mathbf{V}}_\Omega$ complete the proof. \blacksquare

5.2.2 Partial Vandermonde Learning With Errors (PV-LWE)

Similar to the standard Knapsack problem, we can define a dual problem, called Partial Vandermonde Learning With Errors (PV-LWE). Whereas the Partial Vandermonde Knapsack problem is defined with respect to the matrix $\overline{\mathbf{V}}_\Omega$, its transpose $\overline{\mathbf{V}}_\Omega^T$ is used for Partial Vandermonde Learning With Errors. We note that this problem is not needed in the security proof of PASS Encrypt (see Chapter 7) but we think that the duality between PV-Knap and PV-LWE (as we show in the next section) deepens our understanding of PV-Knap (which is used for the security of PASS Encrypt). Furthermore, we think that PV-LWE is an interesting problem in its own, as it may be the starting point of other cryptographic schemes and thus we hope to stimulate further research in this direction. In the following, we use the same notation as in Section 5.2.1.

Definition 5.4 (Partial Vandermonde Learning With Errors): Given n, t, q as above defining the set \mathcal{P}_t . Let ψ be a distribution over \mathbb{Z}^n and fix $\mathbf{s} \in \mathbb{Z}_q^t$. Let $\mathcal{B}_{\mathbf{s}, \psi}$ denote the Partial Vandermonde Learning With Errors distribution over $\mathbb{Z}_q^{t \times n} \times \mathbb{Z}_q^n$, obtained by sampling $\Omega \leftarrow$

$U(\mathcal{P}_t)$, $\mathbf{e} \leftarrow \psi$ and returning $(\overline{\mathbf{V}}_\Omega, \mathbf{b} = \overline{\mathbf{V}}_\Omega^T \cdot \mathbf{s} + \mathbf{e} \bmod q)$. Partial Vandermonde Learning With Errors comes in two variants:

PV-SLWE $_{\overline{\mathbf{V}}_\Omega, \psi}$: Sample $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$. Given a sample of $B_{\mathbf{s}, \psi}$, find \mathbf{s} .

PV-LWE $_{\overline{\mathbf{V}}_\Omega, \psi}$: Sample $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$. Distinguish between a sample from $B_{\mathbf{s}, \psi}$ and a sample of the form $(\overline{\mathbf{V}}_\Omega, \mathbf{b})$, where $\mathbf{b} \leftarrow U(\mathbb{Z}_q^n)$.

An instance of PV-SLWE defines an instance of BDD (Def. 1.4) in the ideal lattice

$$\Lambda_q(\overline{\mathbf{V}}_\Omega) = \{\mathbf{a} \in R: \mathbf{a} = \overline{\mathbf{V}}_\Omega^T \cdot \mathbf{s} \bmod q \text{ for some } \mathbf{s} \in \mathbb{Z}_q^t\},$$

i.e., the lattice generated by the rows of $\overline{\mathbf{V}}_\Omega$ (see Section 1.2.1). Lemma 5.2 below shows that it is not only closed with regard to addition, but also closed with regard to multiplication by any ring element, making it indeed an ideal lattice in R .

Lemma 5.2

Let $\mathbf{a} \in \Lambda_q(\overline{\mathbf{V}}_\Omega)$ and $\mathbf{r} \in R$. Then, $\mathbf{r} \cdot \mathbf{a} \in \Lambda_q(\overline{\mathbf{V}}_\Omega)$.

Proof: Let $\mathbf{a} \in \Lambda_q(\overline{\mathbf{V}}_\Omega)$, i.e., $\mathbf{a} = \overline{\mathbf{V}}_\Omega^T \cdot \mathbf{s} \bmod q$ for some vector $\mathbf{s} \in \mathbb{Z}_q^t$. Note that the multiplication $\mathbf{r} \cdot \mathbf{a}$ is done in $R = \mathbb{Z}[x]/\langle f(x) \rangle$. Let $\text{Rot}(\mathbf{r}) = (r_{\ell k})_{\ell, k \in [n]} \in \mathbb{Z}^{n \times n}$ denote the matrix of multiplication of \mathbf{r} in R with respect to the coefficient embedding, as defined in Section 1.1.3. Then, it yields for $\ell \in [n]$ that $(\mathbf{r} \cdot \mathbf{a})_\ell = (\text{Rot}(\mathbf{r}) \cdot \mathbf{a})_\ell = \sum_{k \in [n]} r_{\ell k} \cdot a_k$. Using that $\mathbf{a} = \overline{\mathbf{V}}_\Omega^T \cdot \mathbf{s} = (\sum_{j \in [t]} s_j \cdot \omega_{ij}^{k-1})_{k \in [n]}$ gives for $\ell \in [n]$ that

$$(\mathbf{r} \cdot \mathbf{a})_\ell = \sum_{k \in [n]} r_{\ell k} \left(\sum_{j \in [t]} s_j \cdot \omega_{ij}^{k-1} \right) = \sum_{j \in [t]} \left[\sum_{k \in [n]} s_j r_{\ell k} \omega_{ij}^{k-1} \right] \omega_{ij}^{\ell-1} = \sum_{j \in [t]} s'_j \omega_{ij}^{\ell-1},$$

where $s'_j = \sum_{k \in [n]} s_j r_{\ell k} \omega_{ij}^{k-1} \in \mathbb{Z}_q$. Thus, $\mathbf{r} \cdot \mathbf{a} = \overline{\mathbf{V}}_\Omega^T \cdot \mathbf{s}'$, where $\mathbf{s}' = (s'_j)_{j \in [t]}$ and finally $\mathbf{r} \cdot \mathbf{a} \in \Lambda_q(\overline{\mathbf{V}}_\Omega)$. ■

In contrast to the standard LWE problem, we do *not* define a variant of the PV-LWE problem in the so-called Hermite Normal Form, where the secret follows the same distribution as the error, since this variant is easy to solve. More precisely, let $\mathbf{b} = \overline{\mathbf{V}}_\Omega^T \cdot \mathbf{s} + \mathbf{e}$. As the first column of $\overline{\mathbf{V}}_\Omega$ is the vector that only contains 1's, the first coefficient of \mathbf{b} is simply the sum of the coefficients of \mathbf{s} plus the first coefficient of \mathbf{e} . If \mathbf{s} and \mathbf{e} are small, so is the first coefficient of \mathbf{b} , making it trivially distinguishable from uniform.

5.2.3 PASS Problem

We now present a leaky variant of dec-PV-Knap that we call the PASS problem, as it is used as the underlying hard problem of PASS Encrypt (from Chapter 7). As opposed to the problems before, it does not only make use of the partial Vandermonde transform $\overline{\mathbf{V}}_\Omega$, but simultaneously

also uses its complement $\overline{\mathbf{V}}_{\Omega^c}$.

Definition 5.5 (PASS Problem): Given n, t, q as above defining the set \mathcal{P}_t . Let ψ_f, ψ_r and ψ_s be distributions over \mathbb{Z}^n . The problem $\text{PASS}_{\psi_f, \psi_r, \psi_s}$ asks to distinguish the following two cases, when given

$$(\overline{\mathbf{V}}_{\Omega} \mathbf{f}, \mathbf{b}, \overline{\mathbf{V}}_{\Omega^c} \mathbf{r}, \overline{\mathbf{V}}_{\Omega^c} \mathbf{s}) \in \left(\mathbb{Z}_q^t \right)^2 \times \left(\mathbb{Z}_q^{n-t} \right)^2.$$

In the first case $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{f} \leftarrow \psi_f$, $\mathbf{r} \leftarrow \psi_r$, $\mathbf{s} \leftarrow \psi_s$ and $\mathbf{b} = \overline{\mathbf{V}}_{\Omega}(\mathbf{f} \cdot \mathbf{r} + \mathbf{s})$. In the second case, the only difference is that $\mathbf{b} \leftarrow U(\mathbb{Z}_q^t)$.

Let ψ denote the distribution of $\mathbf{a} = \mathbf{f} \cdot \mathbf{r} + \mathbf{s}$. Intuitively, the vector \mathbf{a} can be seen as the secret of an instance $(\overline{\mathbf{V}}_{\Omega}, \mathbf{b})$ of dec-PV-Knap $_{\psi}$, where we are given additional information in the form of $\overline{\mathbf{V}}_{\Omega} \mathbf{f}, \overline{\mathbf{V}}_{\Omega^c} \mathbf{r}$ and $\overline{\mathbf{V}}_{\Omega^c} \mathbf{s}$, which we interpret as some leakage on the secret \mathbf{a} . It is clear that this problem is easier than the standard dec-PV-Knap. In Chapter 7 in Section 7.3.3 we provide some refined analysis on the concrete hardness of PASS.

5.3 Equivalence of PV-Knap and PV-LWE

We now show the equivalence of the problems PV-Knap and PV-LWE for the special case of power-of-2 cyclotomics. Let ν be a power of two, $n = \nu/2$ and $t = n/2$. As before, q is set to be a prime such that $q \equiv 1 \pmod{\nu}$. Further, let ω be a primitive ν -th root of unity modulo q and $\gamma = \omega^2$ be a primitive n -th root of unity modulo q . We sample $\Omega \leftarrow U(\mathcal{P}_t)$, and as ν is a power of two, the set Ω only contains odd powers of ω . We denote by $\Omega^c := \{\omega^j : j \in \mathbb{Z}_\nu^\times\} \setminus \Omega$ the complement set, which also only contains odd powers of ω . Additionally, we define $\text{inv}(\overline{\mathbf{V}}_{\Omega}) \in \mathbb{Z}_q^{t \times n}$ as the matrix whose (j, ℓ) -th element is the inverse in \mathbb{Z}_q of the (j, ℓ) -th element of $\overline{\mathbf{V}}_{\Omega}$ for $j \in [t]$ and $\ell \in [n]$. Since $\text{inv}(\overline{\mathbf{V}}_{\Omega})$ is just $\overline{\mathbf{V}}_{\text{inv}(\Omega)}$, where $\text{inv}(\Omega)$ is the set of inverses of the elements of Ω , we write $\overline{\mathbf{V}}_{\text{inv}(\Omega)}$ instead of $\text{inv}(\overline{\mathbf{V}}_{\Omega})$ in the following.

Lemma 5.3 (Duality of Partial Vandermonde Matrices)

Given n, q, Ω as above, defining $\overline{\mathbf{V}}_{\Omega}$ and $\overline{\mathbf{V}}_{\Omega^c}$. It yields

$$\overline{\mathbf{V}}_{\Omega} \cdot \overline{\mathbf{V}}_{\text{inv}(\Omega^c)}^T = \overline{\mathbf{V}}_{\Omega^c} \cdot \overline{\mathbf{V}}_{\text{inv}(\Omega)}^T = \mathbf{0} \in \mathbb{Z}_q^{t \times t}.$$

Proof: We show the first part, the second follows in an analogue manner. For $j, \ell \in [t]$, the entry of the j -th row and the ℓ -th column of the matrix product $\overline{\mathbf{V}}_{\Omega} \cdot \overline{\mathbf{V}}_{\text{inv}(\Omega^c)}^T$ is given by

$$\sum_{k \in [n]} (\omega_j \cdot \omega_\ell^{-1})^{k-1} = \sum_{k \in [n]} (\omega^{i_j} \cdot \omega^{-i_\ell})^{k-1} = \sum_{k \in [n]} (\omega^u)^{k-1} = \sum_{k \in [n]} (\gamma^v)^{k-1},$$

where $\omega_j = \omega^{i_j} \in \Omega$ and $\omega_\ell = \omega^{i_\ell} \in \Omega^c$ and thus $0 \neq i_j - i_\ell = u$, with $u \equiv 0 \pmod{2}$. We can thus write $u = 2v$ for some non-zero integer v and deduce that $\omega^u = \gamma^v$. The last geometric sum $T := \sum_{k \in [n]} (\gamma^v)^{k-1}$ satisfies $(1 - \gamma^v) \cdot T = 1 - (\gamma^v)^n$. As γ is a primitive n -th root of unity and $0 < v < n$, we have that γ^v is an n -th root of unity and $\gamma^v \neq 1$, so the last sum $T = 0$. ■

The proof of Lemma 5.3 uses the special shape of Ω in the case of power-of-2 cyclotomic fields. It is, however, not true for general cyclotomic fields, as illustrated in the following counter example.

🔔 **Example 5.2 (Counter Example for Prime Cyclotomics)**

We consider the case, where ν and q are prime integers with $q \equiv 1 \pmod{\nu}$. In this case, the minimal polynomial is given by $f(x) = x^{\nu-1} + \dots + x + 1$ (i.e. of degree $\varphi(\nu) = \nu - 1$). For concreteness, we choose $\nu = 5$. Thus, the 5-th cyclotomic ring is given by $R = \mathbb{Z}_q[x]/\langle f(x) \rangle$, with $f(x) = x^4 + x^3 + x^2 + x + 1$. Let ω denote a primitive 5-th root of unity in \mathbb{Z}_q , fulfilling $\omega^5 = 1 \pmod{q}$ and $\omega^4 + \omega^3 + \omega^2 + \omega + 1 = 0 \pmod{q}$. A possible choice could be $\omega = 3$ for $q = 11$. Let $\Omega = \{\omega, \omega^2\} \subset \{\omega^j : 1 \leq j \leq 4\}$ and thus $\Omega^c = \{\omega^3, \omega^4\}$. Then, the partial Vandermonde matrix $\bar{\mathbf{V}}_\Omega \in \mathbb{Z}_q^{2 \times 4}$ and its complement $\bar{\mathbf{V}}_{\Omega^c} \in \mathbb{Z}_q^{2 \times 4}$ are given by

$$\bar{\mathbf{V}}_\Omega = \begin{bmatrix} 1 & \omega & \omega^2 & \omega^3 \\ 1 & \omega^2 & \omega^4 & \omega \end{bmatrix} \text{ and } \bar{\mathbf{V}}_{\Omega^c} = \begin{bmatrix} 1 & \omega^3 & \omega & \omega^4 \\ 1 & \omega^4 & \omega^3 & \omega^2 \end{bmatrix}.$$

However, it yields

$$\bar{\mathbf{V}}_\Omega \cdot \bar{\mathbf{V}}_{\text{inv}(\Omega^c)}^T = \begin{bmatrix} -\omega^2 & -\omega^3 \\ -\omega & -\omega^2 \end{bmatrix} \neq \mathbf{0} \in \mathbb{Z}_q^{2 \times 2}.$$

We now show the equivalence of dec-PV-Knap and PV-LWE, both in their decision variant. The equivalence of their search versions follows in an analogue manner. The proof is essentially the same as for standard Knapsack and LWE, see for instance [MM11, Sec. 4.2], combined with the duality Lemma 5.3. In Chapter 7 in Section 7.3.5, we show how to interpret their equivalence in terms of error-correcting codes.

Lemma 5.4 (From PV-Knap to PV-LWE)

Let n be a power of two and $t = n/2$. Further, let ψ denote a distribution over \mathbb{Z}^n . There is a PPT reduction from the problem dec-PV-Knap $_\psi$ to PV-LWE $_\psi$.

Proof: Given an efficient algorithm A for PV-LWE $_\psi$, we build an efficient algorithm B for dec-PV-Knap $_\psi$ that is given an instance $(\bar{\mathbf{V}}_\Omega, \mathbf{b})$ with $\mathbf{b} = \bar{\mathbf{V}}_\Omega \mathbf{a} \pmod{q}$ or $\mathbf{b} \leftarrow U(\mathbb{Z}_q^t)$, where $\mathbf{a} \leftarrow \psi$ and $\Omega \leftarrow U(\mathcal{P}_t)$. Set $\bar{\mathbf{V}}_{\Omega'} = \bar{\mathbf{V}}_{\text{inv}(\Omega^c)}$ and note that if $\Omega \leftarrow U(\mathcal{P}_t)$, then $\text{inv}(\Omega^c)$ also follows the uniform distribution over \mathcal{P}_t , where $\text{inv}(\Omega^c) = \{\omega^{-j} = \omega^{\nu-j} : \omega^j \in \Omega^c\} \subseteq \{\omega^k : k \in \mathbb{Z}_\nu^\times\}$. Thus, $\bar{\mathbf{V}}_{\Omega'}$ defines a valid matrix for PV-LWE. The algorithm B for dec-PV-Knap samples $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$ and picks an arbitrary preimage $\mathbf{v} \in R$ of \mathbf{b} under $\bar{\mathbf{V}}_\Omega$, i.e., $\bar{\mathbf{V}}_\Omega \mathbf{v} = \mathbf{b} \pmod{q}$ (such preimage exists since $\bar{\mathbf{V}}_\Omega$ is of full rank t over \mathbb{Z}_q). The algorithm B then sets $\mathbf{b}' = \bar{\mathbf{V}}_{\Omega'}^T \cdot \mathbf{s} + \mathbf{v}$ and runs A on input $(\bar{\mathbf{V}}_{\Omega'}, \mathbf{b}')$, returning whatever A returns. We now argue that in the real case $\mathbf{b} = \bar{\mathbf{V}}_\Omega \cdot \mathbf{a} \pmod{q}$, then $(\bar{\mathbf{V}}_{\Omega'}, \mathbf{b}')$ is a valid real case instance of PV-LWE $_\psi$. Indeed, we know that $\mathbf{v} = \mathbf{v}' + \mathbf{a}$ for some $\mathbf{v}' \in R$ with $\bar{\mathbf{V}}_\Omega \cdot \mathbf{v}' = \mathbf{0} \pmod{q}$. Using Lemma 5.3, we have $\bar{\mathbf{V}}_\Omega \cdot \bar{\mathbf{V}}_{\Omega'}^T = \mathbf{0} \pmod{q}$ and thus, \mathbf{v}' has to be in the image of $\bar{\mathbf{V}}_{\Omega'}^T$ and hence $\mathbf{v}' = \bar{\mathbf{V}}_{\Omega'}^T \cdot \mathbf{s}'$, for some $\mathbf{s}' \in \mathbb{Z}_q^t$ and finally $\mathbf{b}' = \bar{\mathbf{V}}_{\Omega'}^T \cdot (\mathbf{s} + \mathbf{s}') + \mathbf{a}$, so \mathbf{b}' has the correct real case PV-LWE distribution with secret $\mathbf{s}'' := \mathbf{s} + \mathbf{s}'$ uniformly random in \mathbb{Z}_q^t (thanks to the uniformly random

choice of \mathbf{s}) and error \mathbf{a} sampled from ψ , as required. In the random case $\mathbf{b} \leftarrow U(\mathbb{Z}_q^t)$, \mathbf{b}' is uniformly random in \mathbb{Z}_q^n since \mathbf{v} is in the uniformly random coset of $\Lambda_q^\perp(\overline{\mathbf{V}}_\Omega)$ defined by \mathbf{b} , and \mathbf{b}' is also uniformly random in this coset thanks to the uniformly random \mathbf{s}' . Therefore, the advantage of \mathbf{B} is the same as that of \mathbf{A} . \blacksquare

Lemma 5.5 (From PV-LWE to PV-Knap)

Let n be a power of two and $t = n/2$. Further, let ψ denote a distribution over \mathbb{Z}^n . There is a PPT reduction from the problem PV-LWE $_{\psi}$ to dec-PV-Knap $_{\psi}$.

Proof: Given an efficient algorithm \mathbf{A} for dec-PV-Knap $_{\psi}$, we build an efficient algorithm \mathbf{B} for PV-LWE $_{\psi}$, that is given an instance $(\overline{\mathbf{V}}_\Omega, \mathbf{b})$ with either $\mathbf{b} = \overline{\mathbf{V}}_\Omega^T \cdot \mathbf{s} + \mathbf{e} \bmod q$ or $\mathbf{b} \leftarrow U(\mathbb{Z}_q^n)$, where $\Omega \leftarrow U(\mathcal{P}_t)$, $\mathbf{s} \leftarrow U(\mathbb{Z}_q^t)$ and $\mathbf{e} \leftarrow \psi$. With the same argumentation as above, we define $\overline{\mathbf{V}}_{\Omega'} = \overline{\mathbf{V}}_{\text{inv}(\Omega^c)}$ fulfilling $\overline{\mathbf{V}}_{\Omega'} \cdot \overline{\mathbf{V}}_\Omega^T = \mathbf{0} \bmod q$ and following the uniform distribution over \mathcal{P}_t . The algorithm \mathbf{B} then computes $\mathbf{b}' = \overline{\mathbf{V}}_{\Omega'} \cdot \mathbf{b}$ and runs \mathbf{A} on input $(\overline{\mathbf{V}}_{\Omega'}, \mathbf{b}')$, returning whatever \mathbf{A} returns. We now argue that in the real case $\mathbf{b} = \overline{\mathbf{V}}_\Omega^T \cdot \mathbf{s} + \mathbf{e} \bmod q$, then $(\overline{\mathbf{V}}_{\Omega'}, \mathbf{b}')$ is a valid real case instance of dec-PV-Knap $_{\psi}$. Indeed, using Lemma 5.3, we have $\mathbf{b}' = \overline{\mathbf{V}}_{\Omega'} \cdot (\overline{\mathbf{V}}_\Omega^T \cdot \mathbf{s} + \mathbf{e}) = \overline{\mathbf{V}}_{\Omega'} \cdot \mathbf{e}$, with $\mathbf{e} \leftarrow \psi$. In the random case $\mathbf{b} \leftarrow U(\mathbb{Z}_q^n)$, \mathbf{b}' is uniformly random in \mathbb{Z}_q^t , as the matrix $\overline{\mathbf{V}}_{\Omega'}$ has full rank t . Hence, \mathbf{B} has the same advantage as \mathbf{A} . \blacksquare

5.4 Efficient Computation of Partial Vandermonde Transform

In Chapter 7, we introduce an encryption scheme whose security is based on the hardness of dec-PV-Knap and PASS. This is why it works with polynomials in the finite ring $R_q = \mathbb{Z}_q[x]/\langle x^n + 1 \rangle$, where n is a power of 2. In order to implement this scheme efficiently, it is crucial to select a well-suited encoding for those mathematical objects. As former works on structured variants of SIS and LWE (e.g. [ADPS16]), we choose the Number Theoretic Transformation (NTT) to encode those in quasi-logarithmic time.

In this section, we argue that the NTT can also be used for the partial Vandermonde transform. More precisely, we show that the multiplication of the partial Vandermonde transformation matrix $\overline{\mathbf{V}}_\Omega \in \mathbb{Z}_q^{t \times n}$ by an element $\mathbf{f} \in R$ can be done efficiently by using a t -point NTT in the special case where $K = \mathbb{Q}(\zeta)$ is a power-of-2 cyclotomic field and where $\overline{\mathbf{V}}_\Omega$ and $\overline{\mathbf{V}}_{\Omega^c}$ are set as the odd and even rows of $\overline{\mathbf{V}}$, respectively. Thus, t becomes the important asymptotic parameter. This corresponds to our observations about the concrete security presented later in Section 7.3.

We remark that in this case, we fix the matrix $\overline{\mathbf{V}}_\Omega$ instead of choosing it at random. A designer of a public key encryption scheme may prefer to increase the entropy of their scheme via the randomness over the choice of $\overline{\mathbf{V}}_\Omega$ by selecting at random a subset Ω of size t . By doing so, they have to accept the higher costs of computing a (partial) n -point NTT instead of a t -point NTT.

In the following, we briefly recall the basic definition of the NTT. For more details on the NTT and the description of fast software implementations in quasi-logarithmic time we refer to [GOPS13]. Let \mathbf{f} be a polynomial in $R_q = \mathbb{Z}_q[x]/\langle x^n + 1 \rangle$ and let γ be a n -th primitive root of unity modulo q . For $k \in \{0, \dots, n-1\}$, the k -th coefficient of the n -point NTT of \mathbf{f} is defined

as $\text{NTT}(\mathbf{f})_k = \sum_{j=0}^{n-1} f_j \gamma^{kj}$.⁴ In an analogue manner, for $t = n/2$ and $\xi = \gamma^2$ a t -th primitive root of unity modulo q , and $\mathbf{g} \in \mathbb{Z}_q[x]/\langle x^t + 1 \rangle$, the k -th coefficient of the t -point NTT of \mathbf{g} is defined as $\text{NTT}(\mathbf{g})_k = \sum_{j=0}^{t-1} g_j \xi^{kj}$, where $k \in \{0, \dots, t-1\}$.

5.4.1 Computation of the Partial Vandermonde Transformation

Let ν be a power of 2, $n = \varphi(\nu) = \nu/2$ and $t = n/2$. Further, let $q \in \mathbb{Z}$ be a prime such that $q \equiv 1 \pmod{\nu}$, thus the ν -th cyclotomic polynomial $x^n + 1$ totally splits over \mathbb{Z}_q . Let $\zeta = \exp(-2\pi i/\nu)$ be a complex primitive ν -th root of unity defining the ν -th cyclotomic number field $K = \mathbb{Q}(\zeta) \cong \mathbb{Q}[x]/\langle x^n + 1 \rangle$ with its ring of integers given by $R = \mathbb{Z}[\zeta] \cong \mathbb{Z}[x]/\langle x^n + 1 \rangle$. We further pick a primitive ν -th root of unity modulo q and denote it by ω . We also set $\gamma = \omega^2$ as a primitive n -th root of unity modulo q and $\xi = \gamma^2 = \omega^4$ as a primitive t -th root of unity modulo q .

We define $\Omega = \{\omega^j : j = 4k + 1, k \in \{0, \dots, t-1\}\}$ and $\Omega^c = \{\omega^j : j = 4k + 3, k \in \{0, \dots, t-1\}\}$, both sets of size t and disjoint. The union $\Omega \cup \Omega^c$ is the set of size n of all primitive ν -th roots of unities modulo q . Thus, $\bar{\mathbf{V}}_\Omega$ and $\bar{\mathbf{V}}_{\Omega^c}$ are composed by the odd and even rows of $\bar{\mathbf{V}}$, respectively. For every $k \in \{0, \dots, t-1\}$ it yields

$$(\bar{\mathbf{V}}_\Omega \cdot \mathbf{f})_k = \sum_{j=0}^{n-1} \omega^{(4k+1)j} f_j = \sum_{j=0}^{t-1} \omega^j f_j \omega^{4kj} + \sum_{j=0}^{t-1} \omega^{j+t} f_{j+t} \omega^{4k(j+t)} = \sum_{j=0}^{t-1} (\omega^j f_j + \omega^{j+t} f_{j+t}) \xi^{kj},$$

where we use that $\omega^{4kt} = (\xi^t)^k = 1$ for all $k \in \{0, \dots, t-1\}$, as ξ is a primitive t -th root of unity modulo q . This corresponds to the k -th coefficient of the t -point NTT of \mathbf{g} with $g_j = \omega^j f_j + \omega^{j+t} f_{j+t}$, where we coefficient-wise multiply f_j with powers of ω in a pre-processing step. The pre-processing step requires $2t = n$ multiplications and t additions. In an analogue manner, we can see that for every $k \in \{0, \dots, t-1\}$

$$(\bar{\mathbf{V}}_{\Omega^c} \cdot \mathbf{f})_k = \sum_{j=0}^{n-1} \omega^{(4k+3)j} f_j = \sum_{j=0}^{t-1} \omega^{3j} f_j \omega^{4kj} + \sum_{j=0}^{t-1} \omega^{3j+t} f_{j+t} \omega^{4k(j+t)} = \sum_{j=0}^{t-1} (\omega^{3j} f_j + \omega^{j+t} f_{j+t}) \xi^{kj}.$$

5.4.2 Computation of the Transposed Partial Vandermonde Transformation

In this section we show that the multiplication of the transposed partial Vandermonde transformation matrix $\bar{\mathbf{V}}_\Omega^T \in \mathbb{Z}_q^{n \times t}$ with an element $\mathbf{s} \in \mathbb{Z}_q^t$ can also be done efficiently by using the NTT of \mathbf{s} . Again, we assume that $\bar{\mathbf{V}}_\Omega$ and $\bar{\mathbf{V}}_{\Omega^c}$ are composed of the odd and even rows of $\bar{\mathbf{V}}$, respectively. More precisely for every $k \in \{0, \dots, n\}$

$$\left(\bar{\mathbf{V}}_\Omega^T \cdot \mathbf{s} \right)_k = \sum_{j=0}^{t-1} s_j \omega^{k(4j+1)} = \omega^k \sum_{j=0}^{t-1} s_j \xi^{kj}.$$

For $k \in \{0, \dots, t-1\}$, this obviously corresponds to a t -point NTT, were we do an additional multiplication by ω^k for every coefficient. For $k \in \{t, \dots, n-1\}$, we use again that $\xi^t = 1$ to see that $(\bar{\mathbf{V}}_\Omega^T \cdot \mathbf{s})_k = \omega^t \cdot (\bar{\mathbf{V}}_\Omega^T \cdot \mathbf{s})_{k-t}$. Thus, with the help of t further multiplications, we obtain the last t

⁴The difference of the NTT transform $(\text{NTT}(\mathbf{f})_k)_{k \in \{0, \dots, n-1\}} \in \mathbb{Z}_q^n$ and the discrete Vandermonde transform $\bar{\mathbf{V}} \cdot \mathbf{f} \in \mathbb{Z}_q^n$ is quite subtle. In the first case we sum over the powers $\gamma^0, \dots, \gamma^{n-1}$, where γ is a n -th primitive root of unity, and in the second case we sum over the powers $\omega^0, \dots, \omega^{n-1}$, where ω is a $2n$ -th primitive root of unity.

coefficients of $\bar{\mathbf{V}}_{\Omega}^T \cdot \mathbf{s}$. Again, this also holds for the complement transposed partial Vandermonde transformation matrix $\bar{\mathbf{V}}_{\Omega^c}^T$, as for every $k \in \{0, \dots, n-1\}$ it yields

$$\left(\bar{\mathbf{V}}_{\Omega^c}^T \cdot \mathbf{s}\right)_k = \sum_{j=0}^{t-1} s_j \omega^{k(4j+3)} = \omega^{3k} \sum_{j=0}^{t-1} s_j \xi^{kj}.$$

Part II

Cryptographic Constructions

Chapter 6

Encryption Based on Middle-Product LWR

This chapter can be seen as the continuation of Chapter 4 and is therefore based on the same joint work with Shi Bai, Dipayan Das, Adeline Roux-Langlois, Weiqiang Wen and Zhenfei Zhang which is published in the proceedings of the conference Asiacrypt 2019 [BBD⁺19].

Contents

6.1	Introduction	103
6.1.1	Our Contributions	104
6.1.2	Related Work	104
6.1.3	Roadmap	104
6.2	Reconciliation	104
6.3	Public Key Encryption Based on MP-LWR	106
6.3.1	Correctness	106
6.3.2	Security	107
6.4	Parameters and Security	111
6.4.1	Asymptotic Parameters	111
6.4.2	Concrete Security	112

6.1 Introduction

In Chapter 4, we introduce a new hardness assumption, the problem Middle-Product Computational Learning With Rounding (MP-CLWR, Def. 4.3). It simultaneously benefits from the simplicity of Learning With Rounding (LWR, Def. 1.17) and the security guarantees of Middle-Product Learning With Errors (MP-LWE, Def. 1.15). The *computational* notion, as opposed to the more standard search and decision notions, was introduced by Chen et al. [CZZ18], together with an efficient construction of a PKE scheme with security proof in the Random Oracle Model (ROM). We recall the definition of the ROM in Section 1.5.2. The random oracle is needed to transform the decision problem defined by the IND-CPA security game (Def. 1.20) into a computational problem. To this end, they make use of a reconciliation mechanism, that helps to correctly decrypt a ciphertext.

6.1.1 Our Contributions

In order to show the usefulness of our newly defined hardness assumption MP-CLWR for cryptography, we present in the following a Public Key Encryption (PKE) scheme whose hardness is implied by MP-CLWR. Its design is simultaneously inspired by two existing PKE schemes: on the one hand, the PKE scheme from Roca et al. [RSSS17] whose security is based on MP-LWE; and on the other hand, the PKE scheme from Chen et al. [CZZ18] whose security is based on the hardness of Computational Learning With Rounding Over Rings.

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 the PKE scheme of [RSSS17]. Furthermore, we save in bandwidth, as the second part of the public key is modulo p and not modulo q , where $p \leq q$ by some order of magnitude. In Section 6.4.1, we provide a more detailed comparison between both PKE schemes. Furthermore, we provide in Section 6.4.2 a study of the concrete security of our scheme by looking at the currently best known attacks against it.

6.1.2 Related Work

Since its introduction in 2017, MP-LWE served as the underlying hardness assumption for several cryptographic primitives: from basic encryption [RSSS17] and signature schemes [BDH⁺20] to lattice trapdoors and identity-based encryption [LVV19] and more advanced hierarchical identity-based encryption [LDSP20] and ring signatures [DAZ19]. The public key encryption scheme *Titanium* [SSZ17], which was submitted to the NIST standardization process, is basing its hardness on MP-LWE. Sakzad et al. [SSZ19] further improved practical aspects of PKE schemes based on the MP-LWE problem by providing a tighter security reduction and a more accurate cryptanalysis. It can be seen as an optimization of the originally proposed scheme by Roca et al. [RSSS17]. To the best of our knowledge, we are the first and only ones who constructed an encryption scheme whose security is based on the Middle-Product Learning With Rounding variant. Regarding the standard Learning With Rounding problem, due to its simplicity and efficiency, several encryption schemes such as *SABER* [DKRV18] (third round) and *Round5* [BBF⁺19] (second round) participating in the NIST standardization process are basing their hardness on (structured variants of) decision LWR.

6.1.3 Roadmap

The rest of the chapter is structured as follows. We start with recalling a reconciliation mechanism in Section 6.2, which serves to obtain correctness of our encryption scheme. In Section 6.3, we define the scheme, then prove its correctness in Section 6.3.1 and prove its security based on the hardness of MP-CLWR in the ROM in Section 6.3.2. Finally, in Section 6.4 we provide asymptotic parameters and explain how to analyze the best known attacks against our scheme, which can be used to derive concrete parameters.

6.2 Reconciliation

Reconciliation is a method used by two parties to agree on a secret bit, when they only share a common value up to an approximation factor. As for the PKE scheme of Chen et al. [CZZ18], we need a reconciliation mechanism in order to guarantee the correctness of our scheme.

A first reconciliation mechanism is given by Ding et al. [DXL12] followed by other proposals (e.g., [Pei14, ADPS16]). We use the notation of Peikert and exert the nearest integer rounding. For this purpose, we need the modular rounding function $\lfloor \cdot \rfloor_2 : \mathbb{Z}_q \rightarrow \mathbb{Z}_2$ for $q \in \mathbb{N}$ (see Definition 1.16 for $p = 2$) and define the reconciliation cross-rounding function $\langle \cdot \rangle_2 : \mathbb{Z}_q \rightarrow \mathbb{Z}_2$ as

$$\langle x \rangle_2 = \left\lfloor \left(\frac{4}{q} \right) \cdot x \right\rfloor \bmod 2.$$

For q even, the reconciliation algorithm REC takes as input two values $y \in \mathbb{Z}_q$ and $b \in \{0, 1\}$ and outputs $\lfloor x \rfloor_2$, where x is the closest element to y such that $\langle x \rangle_2 = b$. A concrete definition of REC is given as follows. Define two disjoint intervals $I_0 = \{0, \dots, \lfloor \frac{q}{4} \rfloor - 1\}$ and $I_1 = \{-\lfloor \frac{q}{4} \rfloor, \dots, -1\}$. Let E be the set given by $[-\frac{q}{8}, \frac{q}{8}] \cap \mathbb{Z}$. Further, let y be an element of \mathbb{Z}_q and b be a bit. Then,

$$\text{REC}(y, b) = \begin{cases} 0, & y \in I_b + E \bmod q, \\ 1, & \text{else.} \end{cases}$$

We recall the following results about the cross-rounding function and the reconciliation mechanism from Peikert [Pei14].

Lemma 6.1 (Claim 3.1 and 3.2 [Pei14])

Let $q \in \mathbb{N}$ be even. If $x \in \mathbb{Z}_q$ follows the uniform distribution over \mathbb{Z}_q , then does $\lfloor x \rfloor_2$ follow the uniform distribution over \mathbb{Z}_2 , given $\langle x \rangle_2$. If $x, y \in \mathbb{Z}_q$ such that $|x - y| < \frac{q}{8}$, then $\text{REC}(y, \langle x \rangle_2) = \lfloor x \rfloor_2$.

However, if q is odd, then the output bit of the reconciliation method is biased. In order to mitigate the bias, Peikert [Pei14] introduces a randomized doubling function.

Definition 6.1: Let $q \in \mathbb{N}$ be odd. The doubling function $\text{DBL} : \mathbb{Z}_q \rightarrow \mathbb{Z}_{2q}$ is defined as $\text{DBL}(x) = 2x - e$, where $e \leftarrow \{-1, 0, 1\}$ with probabilities $\Pr[e = -1] = \Pr[e = 1] = \frac{1}{4}$ and $\Pr[e = 0] = \frac{1}{2}$.

Now, analogue results also hold for an odd modulus q .

Lemma 6.2 (Claim 3.3 [Pei14])

Let $q \in \mathbb{N}$ be odd. If $x \in \mathbb{Z}_q$ follows the uniform distribution over \mathbb{Z}_q , then does $\lfloor \bar{x} \rfloor_2$ with $\bar{x} \leftarrow \text{DBL}(x)$ follow the uniform distribution over \mathbb{Z}_2 , given $\langle \bar{x} \rangle_2$. If $x, y \in \mathbb{Z}_q$ such that $|x - y| < \frac{q}{8}$, let $\bar{x} \leftarrow \text{DBL}(x)$, then $\text{REC}(2y, \langle \bar{x} \rangle_2) = \lfloor \bar{x} \rfloor_2$.

We extend all functions $\langle \cdot \rangle_2$, $\lfloor \cdot \rfloor_2$ and $\text{DBL}(\cdot)$ component-wise to vectors over \mathbb{Z}_q and coefficient-wise to polynomials in $\mathbb{Z}_q[x]$, as well as the mechanism REC to vectors over $\mathbb{Z}_q \times \{0, 1\}$ and to polynomials in $\mathbb{Z}_q[x] \times \{0, 1\}[x]$.

6.3 Public Key Encryption Based on MP-LWR

In the following, we define the PKE scheme. We use the rounding functions $\langle \cdot \rangle_2, \lfloor \cdot \rfloor_2$ and the reconciliation method $\text{REC}(\cdot, \cdot)$ from Section 6.2. Without loss of generality, we assume that the modulus q is even. For q odd, we simply need to use the randomized doubling function $\text{DBL}(\cdot)$ as in Section 6.2 to mitigate the bias in the reconciliation mechanism. Further recall from Section 1.4.5 the probabilistic lifting function $\text{lift}(\cdot)$ from $\mathbb{Z}_p[x]$ to $\mathbb{Z}_q[x]$ for two integers p and q with $2 \leq p \leq q$. It lifts rounded polynomials in $\mathbb{Z}_p[x]$ to $\mathbb{Z}_q[x]$ such that $\left\lfloor \text{lift}(\lfloor a \rfloor_p) \right\rfloor_p = \lfloor a \rfloor_p$. Note that $\text{lift}(\lfloor a \rfloor_p) = a + e$ with $\|e\|_\infty \leq q/p$, and we say that lifting introduces a rounding error.

Let k, d, n, p, q and t be positive integers with $d+k \leq n$ and $q \geq p \geq 2$. Further, let H denote a random oracle $H: \{0, 1\}^d \rightarrow \{0, 1\}^k$. The message space is $M = \{0, 1\}^k$ and the security parameter is λ . By \oplus we denote the bit-wise XOR operation of bit strings.

Protocol 6.1: The Public Key Encryption Scheme

KGen(1^λ). Sample $s \leftarrow U((\mathbb{Z}_q^{\leq n+d+k-1}[x])^\times)$,
for $j \in [t]$ choose $a_j \leftarrow U(\mathbb{Z}_q^{\leq n}[x])$ and compute $b_j = \lfloor a_j \odot_{d+k} s \rfloor_p$,
return $\text{pk} = (a_j, b_j)_{j \in [t]}$ and $\text{sk} = s$.

Enc(pk, m). For $j \in [t]$, sample $r_j \leftarrow U(\{0, 1\}^{\leq k+1}[x])$
and set $c_1 = \sum_{j \in [t]} r_j a_j \bmod q$,
compute $v = \sum_{j \in [t]} r_j \odot_d \text{lift}(b_j) \bmod q$
and set $c_2 = \langle v \rangle_2$ and $c_3 = H(\lfloor v \rfloor_2) \oplus m$,
return $c = (c_1, c_2, c_3)$.

Dec(sk, c). Compute $w = c_1 \odot_d s$ and return $m' = c_3 \oplus H(\text{REC}(w, c_2))$.

The encryption scheme, illustrated in Protocol 6.1, works as follows. During key generation, we sample s uniformly at random over $(\mathbb{Z}_q^{\leq n+d+k-1}[x])^\times$ such that $\text{Hank}(s)$ has full rank, defining the secret key $\text{sk} = s$. This can be done by sampling s uniformly at random and rejecting it if its Hankel matrix is not full rank. Lemma 4.1 quantifies the proportion of random s that fulfill this condition. For $j \in [t]$, we sample polynomials $a_j \in \mathbb{Z}_q^{\leq n}[x]$ uniformly at random and compute the rounded middle-product $b_j = \lfloor a_j \odot_{d+k} s \rfloor_p$, defining the public key $\text{pk} = (a_j, b_j)_{j \in [t]}$. In order to encrypt a message $m \in \{0, 1\}^k = M$, we sample for $j \in [t]$ a polynomial $r_j \in \{0, 1\}^{\leq k+1}[x]$ uniformly at random and set $c_1 = \sum_{j \in [t]} r_j a_j \bmod q$, defining the first part of the ciphertext. Here we use within the sum the standard convolution product of two polynomials. Then, we compute an intermediate value $v = \sum_{j \in [t]} r_j \odot_d \text{lift}(b_j) \bmod q$, which is used to compute the second and third ciphertext part as follows: set $c_2 = \langle v \rangle_2$ and $c_3 = H(\lfloor v \rfloor_2) \oplus m$. Note, that this time we use within the sum the middle-product of two polynomials. Finally, output the ciphertext $c = (c_1, c_2, c_3) \in \mathbb{Z}_q^{\leq n+k}[x] \times \{0, 1\}^d \times \{0, 1\}^k$. In order to decrypt a ciphertext c for a secret key s we first compute the intermediate value $w = c_1 \odot_d s$ and then return $m' = c_3 \oplus H(\text{REC}(w, c_2))$.

6.3.1 Correctness

We now show that the PKE scheme defined above is perfectly correct under a proper choice of parameters. See Definition 1.19 for the formal statement of the correctness property of a PKE scheme.

Lemma 6.3 (Correctness)

Let $t, k, p \in \mathbb{N}$ be part of the public parameters of the encryption scheme from above. Assume that $p > 8t(k+1)$. For every plaintext $m \in M$ and key pair $(pk, sk) \leftarrow \text{KGen}(1^\lambda)$, we have

$$\Pr[\text{Dec}(sk, \text{Enc}(pk, m)) = m] = 1.$$

Proof: In order to prove the correctness of the scheme, we need to guarantee that the reconciliation mechanism succeeds. Actually, this step is independent of the message. Following Lemma 6.1 (and for q odd Lemma 6.2) it suffices to show that $\|w - v\|_\infty < q/8$. It yields

$$\begin{aligned} v &= \sum_{j \in [t]} r_j \odot_d \text{lift}(b_j) = \sum_{j \in [t]} r_j \odot_d (a_j \odot_{d+k} s + e_j) = \sum_{j \in [t]} (r_j a_j) \odot_d s + \sum_{j \in [t]} r_j \odot_d e_j \\ &= c_1 \odot_d s + \sum_{j \in [t]} r_j \odot_d e_j = w + \sum_{j \in [t]} r_j \odot_d e_j, \end{aligned}$$

where for $j \in [t]$ the error term e_j with $\|e_j\|_\infty < q/p$ is determined by the lifting function $\text{lift}(\cdot)$. In the third equation we used the associativity law of the middle-product Lemma 1.24. Thus it suffices to have

$$\left\| \sum_{j \in [t]} r_j \odot_d e_j \right\|_\infty < q/8.$$

For $j \in [t]$ each coefficient of $r_j \odot_d e_j$ can be seen as the inner product $\langle u, v \rangle$ of a binary vector u of dimension $k+1$ and a vector v also of dimension $k+1$, where each coefficient has magnitude $\leq q/p$. Notice that we have

$$|\langle u, v \rangle| \leq \|u\|_2 \cdot \|v\|_2 \leq \sqrt{k+1} \cdot \sqrt{(k+1) \cdot q^2/p^2} = (k+1)q/p.$$

Hence, it yields

$$\left\| \sum_{j \in [t]} r_j \odot_d e_j \right\|_\infty \leq \sum_{j \in [t]} \|r_j \odot_d e_j\|_\infty \leq t(k+1)q/p.$$

As $p > 8t(k+1)$, we have $t(k+1)q/p < q/8$ which guarantees that the reconciliation mechanism succeeds. ■

6.3.2 Security

In this section, we prove the security of the PKE scheme defined above based on the hardness of MP-CLWR in the ROM. We use the standard notion of IND-CPA security whose proper definition we recall in Definition 1.20.

The high level idea of the security proof is the following. We first define a game, which we call REC-COMP, where the adversary has to compute the output of the reconciliation mechanism by only having access to some related information. We then embed the IND-CPA security game into the REC-COMP game. In other words, the adversary of the latter is at the same time the challenger of the first. To this end, we use that a message is encrypted by evaluating a random oracle H on exactly this value and bit-wise XOR-ing it with the message. An adversary who can distinguish the two encrypted messages from the IND-CPA game has to query the random oracle H

on the corresponding input at some point. And thus, the challenger of the IND-CPA game, who also simulates the random oracle, can use this query to solve the REC-COMP challenge. That is to say, the random oracle helps to transform a distinguishing game (the IND-CPA security game) into a computational game (the REC-COMP game).

Before stating and proving the security of our scheme, we prove that the following hash function family using the middle-product is universal. Recall the Definition 1.7 of universal families of hash functions in Section 1.3.5. Further recall that $\text{lift}(\cdot)$ denotes the probabilistic lifting function from $\mathbb{Z}_p[x]$ to $\mathbb{Z}_q[x]$ for two integers p and q with $2 \leq p \leq q$ (see Section 1.4.5).

Lemma 6.4 (Universal Family of Hash Functions)

Let $q, k, d, p, t \in \mathbb{N}$ such that $k, d \geq 2$ and $2 \leq p \leq q$. For $(b_j)_{j \in [t]} \in (\mathbb{Z}_p^{\leq d+k}[x])^t$, we define

$$h_{(b_j)} : \left(\{0, 1\}^{\leq k+1}[x] \right)^t \rightarrow \mathbb{Z}_q^{\leq d}[x], \quad \text{where } (r_j)_j \mapsto \sum_{j \in [t]} \text{lift}(b_j) \odot_d r_j.$$

The hash function family $(h_{(b_j)})_{(b_j)_j}$ is universal.

Proof: The proof is very similar to the one of [RSSS17, Lemma 4.2]. We simply replace b_j by $\text{lift}(b_j)$, using the same argument to show that

$$\Pr_{(b_j)_j \leftarrow U((\mathbb{Z}_p^{\leq d+k}[x])^t)} \left[\sum_{j \in [t]} \text{lift}(b_j) \odot_d r_j = \sum_{j \in [t]} \text{lift}(b_j) \odot_d r'_j \right] = \frac{1}{q^d},$$

with $(r_j)_j \neq (r'_j)_j$. ■

We can now prove the security of our scheme.

Lemma 6.5 (Security)

Let λ be the security parameter and $k, d, n, p, q, t \in \mathbb{N}$ such that $d + k \leq n$ and $q \geq p \geq 2$. Assume that $t \geq (2 \cdot \lambda + (k + d + n) \cdot \log_2 q) / (k + 1)$. The PKE scheme above is IND-CPA secure in the ROM under the $\text{MP-CLWR}_{p,q,n,d+k,t}$ hardness assumption.

Proof: *Setup.* Recall that the IND-CPA security game (Protocol 1.1) is the following: A challenger \mathcal{C} generates a key pair (pk, sk) , samples a random bit b and sends the public key pk to the adversary \mathcal{A} . The adversary chooses two messages m_0, m_1 and sends them to the challenger \mathcal{C} , which in turn encrypts m_b and sends the corresponding ciphertext c back to \mathcal{A} . The adversary outputs a bit b' as a guess of b and wins the game if $b = b'$.

Let $c = (c_1, c_2, c_3)$ be the ciphertext of the message $m_b \in M$. The only part that depends on the message m_b is the last one, where $c_3 = H(\lfloor v \rfloor_2) \oplus m_b$. If the random oracle H was not queried on the value of $\lfloor v \rfloor_2 \in \{0, 1\}^d$ during the game, the adversary \mathcal{A} can only guess the randomly chosen bit b with success probability $1/2$. Thus, we can assume that a successful adversary (with success probability non-negligibly larger than $1/2$) has queried H on this spe-

cific value. Subsequently, we can use a successful adversary \mathcal{A} of the IND-CPA security game to construct a successful adversary \mathcal{A}' which outputs $\lfloor v \rfloor_2$, given the first two parts (c_1, c_2) of any ciphertext $c = (c_1, c_2, c_3)$. These first two parts are independent of the message to encrypt. We call this the REC-COMP game, as it yields $c_2 = \langle v \rangle_2$ and thus c_1 and c_2 are connected via the reconciliation mechanism REC. In Protocol 6.2 we illustrate how the IND-CPA security game is nested into the REC-COMP game. The adversary \mathcal{A}' in the latter plays simultaneously the role of the challenger in the former. Note that the REC-COMP game asks the adversary to *compute* something, whereas the IND-CPA game asks the adversary to *distinguish* something. By using the random oracle H (simulated by \mathcal{A}') to encrypt a message, we can use a successful *distinguisher* in order to obtain a successful *computer*.

During the IND-CPA game, \mathcal{A}' answers the random oracle queries of \mathcal{A} by maintaining an input-output table for H . For each query, \mathcal{A}' first checks if H was already programmed on the queried input. If yes, they output the corresponding hash value, otherwise they choose a fresh random value and set H accordingly. Assuming \mathcal{A} has non-negligible advantage to win the IND-CPA security game, it must have queried H on $\lfloor v \rfloor_2$, hence \mathcal{A}' can look up the pair $(\lfloor v \rfloor_2, r)$ with $r = H(\lfloor v \rfloor_2)$ in the random oracle table.

Game-based proof. As a next step, we show that the probability of \mathcal{A}' to win is negligible under the MP-CLWR assumption, which in turns implies the IND-CPA security of our encryption scheme. We consider the following sequence of games, where in all games $a_j \leftarrow U(\mathbb{Z}_q^{<n}[x])$ for $j \in [t]$ and the secret s is chosen via $s \leftarrow U(\left(\mathbb{Z}_q^{<n+d+k-1}[x]\right)^\times)$. Further, we sample $r_j \leftarrow U(\{0, 1\}^{<k+1}[x])$ for $j \in [t]$ and set the first part of the ciphertext as $c_1 = \sum_{j \in [t]} r_j a_j \bmod q$.

The adversary \mathcal{A}' receives in each game the tuple $(1^\lambda, \text{pk}, c_1, c_2)$ and its target is to compute $\lfloor v \rfloor_2$, where v is specified by each game separately.

Game 1: Set $b_j = \lfloor a_j \odot_{d+k} s \rfloor_p$, $\text{pk} = (a_j, b_j)_j$, $v = \sum_j \text{lift}(b_j) \odot_d r_j \bmod q$, and $c_2 = \langle v \rangle_2$,

Game 2: Set $b_j \leftarrow \left[U(\mathbb{Z}_q^{<d+k}[x]) \right]_p$, $\text{pk} = (a_j, b_j)_j$, $v = \sum_j \text{lift}(b_j) \odot_d r_j \bmod q$, and $c_2 = \langle v \rangle_2$,

Game 3: Set $b_j \leftarrow \left[U(\mathbb{Z}_q^{<d+k}[x]) \right]_p$, $\text{pk} = (a_j, b_j)_j$, $v \leftarrow U(\mathbb{Z}_q^{<d}[x])$, and $c_2 = \langle v \rangle_2$.

Third game. Note that in the last game, c_1 and c_2 are independent and hence the probability that \mathcal{A}' outputs $\lfloor v \rfloor_2 \in \{0, 1\}^d$ is exactly $1/2^d$, using Lemma 6.1 (and Lemma 6.2 for q odd).

From third to second game. Furthermore, the second and third game are within exponentially small statistical distance, using the Generalized Leftover Hash Lemma from Section 1.3.5. In more details, the statistical distance of the two distributions of $((a_j, b_j)_j, c_1, v)$ in Game 2 and Game 3 is given by

$$\Delta \left[\left((a_j, b_j)_j, \sum_{j \in [t]} r_j a_j, \sum_{j \in [t]} r_j \odot_d \text{lift}(b_j) \right), \left((a_j, b_j)_j, \sum_{j \in [t]} r_j a_j, v \right) \right] \leq 2^{-\lambda},$$

where for all $j \in [t]$ the polynomials a_j, b_j, r_j and v are chosen uniformly at random in $\mathbb{Z}_q^{<n}[x]$, $\left[\mathbb{Z}_q^{<d+k}[x] \right]_p$, $\{0, 1\}^{<k+1}[x]$ and $\mathbb{Z}_q^{<d}[x]$, respectively. Note that the randomness of the family of hash functions $(h_{(b_j)_j})_{(b_j)_j}$ comes from the randomness of $(b_j)_j$ and since Lemma 6.4

shows that $(h_{(b_j)_i})_{(b_j)_j}$ is universal we can use Lemma 1.22. Thus, using the assumptions made on the parameters, the statistical distance is bounded above by

$$\frac{1}{2} \cdot \sqrt{2^{-(k+1)t} \cdot q^{k+n+d}} \leq 2^{-\lambda}.$$

From second to first game. The first and second game differ only in the way how the b_j are computed. In the first game, b_j is a rounded middle-product sample and in the latter one, it is a rounded uniform sample. We can interpret this situation as two different experiments, see Protocol 6.3.

Recall from Definition 4.3 that \mathcal{X}_s^t denotes the distribution of $(a_j, \lfloor a_j \odot_d s \rfloor_p)_{j \in [t]}$, where we choose the $a_j \leftarrow U(\mathbb{Z}_q^{\leq n}[x])$ independently and sample a fixed secret element $s \leftarrow U((\mathbb{Z}_q^{\leq n+d+k-1}[x])^\times)$. Further, we denote by \mathcal{U}^t the distribution of $(a_j, \lfloor b_j \rfloor_p)_{j \in [t]}$, where we choose the $a_j \leftarrow U(\mathbb{Z}_q^{\leq n}[x])$ and the $b_j \leftarrow U(\mathbb{Z}_q^{\leq d+k}[x])$ independently. In addition, con is an arbitrary distribution over $\{0, 1\}^*$ which is independent from \mathcal{X}_s^t and \mathcal{U}^t . The Input_1 of the first experiment $\text{Exp}_1(\mathcal{C}, \mathcal{A}, \mathcal{S}_1)$ is given by $(1^\lambda, \text{pk}, c_1, \langle v \rangle_2)$, where $v = \sum_j \text{lift}(b_j) \odot_d r_j$ with $b_j = \lfloor a_j \odot_{d+k} s \rfloor_p$. On the other hand, the Input_2 of the second experiment $\text{Exp}_2(\mathcal{C}, \mathcal{A}, \mathcal{S}_2)$ is defined by $(1^\lambda, \text{pk}, c_1, \langle v \rangle_2)$, where we still have $v = \sum_j \text{lift}(b_j) \odot_d r_j$ but this time with $b_j \leftarrow \left[U(\mathbb{Z}_q^{\leq d+k}[x]) \right]_p$. The Target is in both cases the same, namely $\lfloor v \rfloor_2$.

According to the MP-CLWR assumption, if the success probability for any \mathcal{A} to output the requested $\lfloor v \rfloor_2$ is negligible when $b_j \leftarrow \left[U(\mathbb{Z}_q^{\leq d+k}[x]) \right]_p$, it is also negligible when b_j is an MP-LWR instance.

First game. Note that Game 1 corresponds to the REC-COMP game above. Combining the arguments above shows that the success probability of \mathcal{A}' is negligible under the MP-CLWR assumption, completing the security proof of our PKE scheme. ■

Protocol 6.2: Nesting the IND-CPA security game into the REC-COMP game

$\mathcal{C}'_{\text{REC-COMP}}$	$\mathcal{A}'_{\text{REC-COMP}/\mathcal{C}_{\text{IND-CPA}}}$	$\mathcal{A}_{\text{IND-CPA}}$
$(\text{pk}, \text{sk}) \leftarrow \text{KGen}(1^\lambda)$		
$(c_1, c_2, *) \leftarrow \text{Enc}(\text{pk}, *)$		
	$\xrightarrow{1^\lambda, \text{pk}, c_1, c_2}$	$\xrightarrow{1^\lambda, \text{pk}}$
	$b \leftarrow \{0, 1\}$	$m_0, m_1 \leftarrow \{0, 1\}^k$
	$r \leftarrow \{0, 1\}^k$	$\xleftarrow{m_0, m_1}$
		$\xrightarrow{c=(c_1, c_2, r \oplus m_b)}$
	look up $\lfloor v \rfloor_2$ in H	$\xleftarrow{b'}$
$\xleftarrow{\lfloor v \rfloor_2}$	such that $H(\lfloor v \rfloor_2) = r$	

Protocol 6.3: Experiment setting of the security proof

$\text{Exp}_1(\mathcal{C}, \mathcal{A}, \mathcal{S}_1)$	$\text{Exp}_2(\mathcal{C}, \mathcal{A}, \mathcal{S}_2)$
1: $((a_j, \lfloor a_j \odot_d s \rfloor_p)_j, \text{aux}) \leftarrow (\mathcal{X}_s^t, \text{con})$	1: $((a_j, b_j)_j, \text{aux}) \leftarrow (\mathcal{U}^t, \text{con})$
2: $(\text{Input}_1, \lfloor v \rfloor_2) \leftarrow \mathcal{C}((a_j, b_j)_j, \text{aux})$	2: $(\text{Input}_2, \lfloor v \rfloor_2) \leftarrow \mathcal{C}((a_j, b_j)_j, \text{aux})$
3: $\text{Output}_1 \leftarrow \mathcal{A}(\text{Input}_1)$	3: $\text{Output}_2 \leftarrow \mathcal{A}(\text{Input}_2)$
4: return $\text{Output}_1 = \lfloor v \rfloor_2$	4: return $\text{Output}_2 = \lfloor v \rfloor_2$

6.4 Parameters and Security

In this section, we first provide some asymptotic parameters and compare them with the ones of the PKE scheme from Roşca et al. [RSSS17]. Then, we explain how to analyze the best known attacks against our scheme, which can be used to derive concrete parameters.

6.4.1 Asymptotic Parameters

In the following, we present asymptotic parameters for our PKE scheme from Section 6.3 for a given security parameter λ . Reasonable values for λ are for instance 128 or 256. Our PKE scheme is defined by the parameters k, d, n, p, q and t . We set the dimension $n \geq \lambda$, the middle-product parameters as $k = d = n/2$, the number of samples as $t = \Theta(\log_2(n))$, the general modulus as $q = \Theta(n^{4+c} \log_2(n)^2)$ and the rounding modulus as $p = \Theta(n \log_2(n))$, where $c \in \mathbb{R}$ is an arbitrary positive constant. Here, the symbols O, Ω and Θ denote the standard asymptotic notions of bounded above, below and from both directions, respectively.

Using these parameters, the scheme is perfectly correct (Lemma 6.3) and secure under the MP-CLWR $_{p,q,n,d+k,t}$ assumption (Lemma 6.5). Using Theorem 4.1, this allows us to rely on the MP-LWE $_{q,n,d+k,\psi}$ problem, where the error distribution ψ is B -bounded with $B = O(n^{2+c})$. The distribution ψ is set to be a discrete Gaussian distribution $D_{\alpha q}$ with $\alpha = \Theta(1/(n\sqrt{\log_2(n)}))$. Using the P-LWE $_{q,D_{\beta q}}^f$ to MP-LWE $_{q,n,d+k,D_{\alpha q}}$ reduction from [RSSS17], this in turn prevents attack as [AG11], where $\beta = \Omega(\sqrt{n}/q)$ for any $f(x) \in \mathbb{Z}[x]$ monic of degree n with constant coefficient coprime with q and expansion factor at least n^c .

We now compare the asymptotic parameters of our encryption scheme with the ones given in [RSSS17]. Table 6.1 shows the asymptotic parameters, key sizes and ciphertext sizes for both schemes. The most important parameter is the value $\log_2(q)$ as it dominates the key and ciphertext sizes of both schemes. Asymptotically, in both cases this value is $\Theta(\log_2(n))$, even though the q is asymptotically larger in our scheme.

In general, the sampling cost is one of the intense operations of an encryption scheme. In the encryption scheme of [RSSS17], we need $2 \cdot t + 1$ sampling subroutines, including t from a rounded Gaussian distribution, during key generation and t sampling subroutines during encryption. In contrast, in our case we only need $t + 1$ sampling subroutine during key generation and t sampling subroutines during encryption. Additionally, in our case all sampling is performed over some uniform distribution which is simpler than Gaussian type sampling.

Further, in our encryption scheme we don't need to restrict the modulus q to be prime. Unlike [RSSS17], it works for all integer moduli which are sufficiently large. This gives an advantage on the choice of parameters.

Table 6.1: Comparison of asymptotic parameters, key sizes and ciphertext sizes

Parameter	[RSSS17]	Our work [BBD ⁺ 19]
n	$\geq \lambda$	$\geq \lambda$
c	> 0	> 0
k	$n/2$	$n/2$
d	$n/2$	$n/2$
t	$\Theta(\log_2(n))$	$\Theta(\log_2(n))$
q	$\Theta(n^{2.5+c} \sqrt{\log_2(n)})$	$\Theta(n^{4+c} \log_2(n)^2)$
$\log_2(q)$	$\Theta(\log_2(n))$	$\Theta(\log_2(n))$
α	$\Theta\left(\frac{1}{n\sqrt{\log_2(n)}}\right)$	-
p	-	$\Theta(n \log_2(n))$
B	-	$O(n^{2+c})$
Key size		
sk	$(n + d + k - 1) \cdot \log_2(q)$	$(n + d + k - 1) \cdot \log_2(q)$
pk	$t \cdot ((n + d + k) \log_2(q))$	$t \cdot (n \log_2(q) + (d + k) \log_2(p))$
Ciphertext size		
c_1	$(n + k) \log_2(q)$	$(n + k) \log_2(q)$
c_2	$d \log_2(q)$	d
c_3	-	k

However, the PKE scheme of [RSSS17] is provably secure in the standard model, whereas our proof is in the ROM.

6.4.2 Concrete Security

Unfortunately, neither the security proof in Lemma 6.5 nor the asymptotic parameters from above give guidance on the choice of *concrete* parameters and on the concrete security they provide for the encryption scheme. Parameter derivation is indeed an active research topic for lattice-based cryptography, for both the construction of cryptographic protocols and cryptanalysis, e.g., [APS15, ADPS16, ACD⁺18]. As explained in Section 1.4.6, in practice it is common to derive the concrete parameters for LWE-based schemes by looking at the cost of the best known attacks, such as BKZ with quantum sieving, in the so-called *core* SVP model, see for instance [ADPS16]. Note that those derived parameters don't necessarily fulfill the conditions of the security proof of Lemma 6.5.

In the following, we highlight two approaches for our PKE scheme. The first one interprets the public key as an instance of LWE by proceeding as follows: The public key of our scheme provides an instance of MP-LWR, which is in a first step interpreted as an instance of MP-LWE and in a second step analyzed as an instance of LWE. The second approach is a dual lattice attack on the public key, that uses the weak associativity law of the middle-product.

Public Key as LWE Instance

Step 1: From MP-LWR to MP-LWE. In our scheme from Section 6.3, the public key is given as $\text{pk} = (a_j, b_j)_{j \in [t]}$, where $b_j = \lfloor a_j \odot_{d+k} s \rfloor_p$ for $a_j \leftarrow U(\mathbb{Z}_q^{<n}[x])$ and $s \leftarrow U((\mathbb{Z}_q^{<n+d+k-1}[x])^\times)$ for $j \in [t]$. In order to interpret this MP-LWR instance as an MP-LWE instance, it suffices to lift the second part from $\mathbb{Z}_p^{<d+k}[x]$ to $\mathbb{R}_q^{<d+k}[x]$ by multiplying it with q/p . This introduces a so-called rounding noise, as explained in Section 1.4.5. Hence, the public key defines the following

equation

$$q/p \cdot b_j = a_j \odot_{d+k} s + e_j, \quad (6.1)$$

where e_j is an element of $\mathbb{R}_q^{<d+k}[x]$ with $\|e_j\|_\infty \leq q/p$.

Step 2: From MP-LWE to LWE. We analyze the hardness of solving an MP-LWE instance by interpreting it as an LWE instance, as for today, the best known attacks don't make use of the underlying structure. This is also true for other structured variants such as P-LWE, R-LWE or M-LWE.

In Section 1.4.4, we saw that the middle-product of two polynomials can be expressed in terms of a matrix-vector product using special product matrices \mathbf{T} , see Definition 1.14 and Lemma 1.25. Thus, given Equation 6.1 from above, it yields

$$q/p \cdot \text{rev}(\mathbf{b}_j) = \mathbf{T}^{d+k,n}(a_j) \cdot \text{rev}(\mathbf{s}) + \mathbf{e}_j, \quad (6.2)$$

where $\text{rev}(\mathbf{b}_j)$ and $\text{rev}(\mathbf{s})$ are the coefficient vectors of b_j and s in reverse order. By ignoring the structure of $\mathbf{T}^{d+k,n}$ and interpreting it as random matrix $\mathbf{A}_j \in \mathbb{Z}_q^{(d+k) \times (n+d+k-1)}$ this gives the final LWE equation

$$\tilde{\mathbf{b}}_j = \mathbf{A}_j \cdot \tilde{\mathbf{s}} + \mathbf{e}_j, \quad (6.3)$$

where $\tilde{\mathbf{b}}_j = q/p \cdot \text{rev}(\mathbf{b}_j)$ and $\tilde{\mathbf{s}} = \text{rev}(\mathbf{s})$. Globally speaking, to analyze an instance of LWE, there are two type of attacks: The primal and the dual attack. The high level idea of the primal attack is to recover the LWE secret, exploiting that it is part of an unusually short vector in a given lattice specified by the LWE instance. In contrast, the attacker can use the dual attack to first find a short vector in the dual lattice and then use this short vector to decide whether a given instance is an LWE instance or drawn from the uniform distribution.

To conclude on the concrete security estimations, one can use the LWE Estimator [APS15] or the more recent Leaky LWE Estimator [DDGR20] by plugging in some concrete values for q, n, d, k, p and t and gradually adapt them to obtain the aimed security level. Note that the requirements of Lemma 6.3 should be met in order to guarantee the correctness of the scheme.

Dual Attack against Public Key

As opposed to the previous approach, where the attacker tries to recover the secret key from the public key via interpreting it as an instance of LWE, the attacker can directly use the dual attack to first find a short vector in the dual lattice and then use this short vector to decide whether a given instance is an MP-LWR instance or drawn from the uniform distribution. Note that this technique attacks the *decision* problem, which is not directly used in our scheme. However, in order to deepen our understanding of the scheme's security, it is interesting to study this attack as well.

Recall that the public key $(a_j, b_j)_{j \in [t]}$ defines the Equation 6.1. Now, we consider the following lattice

$$\Lambda_q(a_1, \dots, a_t) = \left\{ (x_1, x_2, \dots, x_t)^T \in \left(\mathbb{Z}^{<k+1}[x] \right)^t : \sum_{j \in [t]} x_j \cdot a_j = 0 \pmod{q} \right\}.$$

If the adversary is able to find a short lattice vector $\mathbf{x} = (x_1, \dots, x_t)^T$ in $\Lambda_q(a_1, \dots, a_t)$, they can

use it to distinguish uniform samples from MP-LWR samples by computing the sum

$$\begin{aligned}
 \sum_{j \in [t]} x_j \odot_d (q/p \cdot b_j) &= \sum_{j \in [t]} x_j \odot_d (a_j \odot_{d+k} s + e_j) \\
 &= \sum_{j \in [t]} (x_j \cdot a_j) \odot_d s + \sum_{j \in [t]} x_j \odot_d e_j \\
 &= \sum_{j \in [t]} x_j \odot_d e_j.
 \end{aligned}$$

In the second equation we use the weak associativity of the middle-product from Lemma 1.24. As the x_j and the e_j are all polynomials with small coefficients, the sum of their product $\sum_j x_j \odot_d e_j$ has (relatively) small coefficients as well. On the other hand, if b_j are uniform elements, the sum $\sum_j x_j \odot_d e_j$ is a polynomial with uniform coefficients, and thus the adversary can distinguish the public key from uniform samples. We refer to [MR09] for the details of the concrete parameter setting for which this attack works.

Concrete Parameters

In the original paper published in the proceedings of Asiacrypt 2019 [BBD⁺19], we also provide concrete sample parameters for our scheme (and for the scheme of Roşca et al. [RSSS17]). As this part of the research was mainly done by the co-author Zhenfei Zhang, we refer to the given parameters there. We would like to notice that those parameters may have become outdated as the work by Sakzad et al. [SSZ19] provides tighter cryptanalysis for schemes based on MP-LWE which most certainly carries over to the rounding setting. As already stated in the original article, we leave a more efficient instantiation and dedicated cryptanalysis to future work.

Chapter 7

PASS Encrypt

This chapter can be seen as the continuation of Chapter 5 and is therefore based on the same joint work with Amin Sakzad and Ron Steinfeld, which was initiated during a research stay at the Monash University in Melbourne from October to December 2019.

Contents

7.1	Introduction	115
7.1.1	Our Contributions	116
7.1.2	Related Work	116
7.1.3	Roadmap	117
7.2	PASS Encrypt	117
7.2.1	Correctness	118
7.2.2	Security	119
7.2.3	Homomorphic Properties	120
7.3	Security Analysis	121
7.3.1	Key Recovery Attack	122
7.3.2	Randomness Recovery Attack	123
7.3.3	Plaintext Recovery Using Hints Attack	123
7.3.4	Choice of Ring	124
7.3.5	Code-Based Attack	125
7.4	Concrete Parameters	125
7.4.1	Choice of the Number of Rows	126
7.4.2	Comparison	127

7.1 Introduction

In Chapter 5, we study several problems related to the discrete Vandermonde matrix $\bar{\mathbf{V}}$. In particular, we define search and decision Partial Vandermonde Knapsack (PV-Knap, Def. 5.2) and the PASS problem (Def. 5.5), where the latter one can be seen as a leaky variant of the first. The motivation of studying those problems is that they serve as the underlying hardness assumptions for the security proof of a Public Key Encryption (PKE) scheme, as we elaborate in the following. Recall from the introduction to Chapter 5 that Hoffstein and Silverman [HS15] use those computational building blocks to construct PASS Encrypt, an efficient lattice-based PKE scheme. On

the positive side, the algebraic structure of the Vandermonde matrix equips PASS Encrypt with additive and multiplicative homomorphic properties. On the negative side, the scheme comes without a security proof with respect to the hardness of explicit computational problems, and the scheme is deterministic and thus cannot satisfy the standard notion of IND-CPA security.¹

7.1.1 Our Contributions

In this chapter, we present a modification of PASS Encrypt together with a security proof based on the decision PV-Knap problem and a leaky variant of it, that we call the PASS problem, both studied in Chapter 5. The latter problem captures the fact that a ciphertext of PASS Encrypt consists of several partial Vandermonde transforms of *related* elements. In other words, a successful attacker against PV-Knap can be used to win the IND-CPA security game, but a successful attacker against the IND-CPA security of PASS Encrypt may not be powerful enough to solve PV-Knap. This issue was not addressed before in the original version of PASS Encrypt [HS15]. Furthermore, the original scheme is deterministic and thus cannot be IND-CPA secure. Additionally, it uses the Fourier transform similar to older versions of PASS Sign.²

In our slightly modified version of PASS Encrypt, we first move from the Fourier to the Vandermonde transform, as done for PASS Sign by Lu et al. [LZA18]. The Fourier transform is defined by the powers of *all* roots of unity, whereas the Vandermonde transform only by the powers of all *primitive* roots of unity. This change is motivated by the fact that the discrete Fourier transform always maps the all-1 vector to zero and thus partial Fourier SIS³ is trivially easy. In contrast, our setting does not allow the same trivial solution to partial Vandermonde SIS. For completeness, we explain how to solve partial Fourier SIS in Section 7.3.4. Second, we make the scheme probabilistic by adding randomness to the message. We then give a proof of correctness (Lemma 7.1) for well-chosen parameters and a proof of security (Lemma 7.2), assuming the hardness of dec-PV-Knap and PASS. A refined analysis of the security of PASS Encrypt is provided in Section 7.3. In particular, we show a novel attack that we call Plaintext Recovering Using Hints attack, which takes the structure of PASS Encrypt into account. It is inspired by the recent work of Dachman-Soled et al. [DDGR20] on exploiting hints that are given on a LWE secret or noise. Our complexity estimates for this attack show that it does not reduce the attack complexity below that of previously known lattice algorithms on PASS Encrypt, which increases our confidence in its claimed security against best known lattice attacks. We conclude the chapter by providing concrete sample parameters and compare our version of PASS Encrypt with two other efficient schemes whose security proofs are based on structured lattice problems.

7.1.2 Related Work

To the best of our knowledge, PASS Encrypt is the only encryption scheme that uses the partial Vandermonde transformation matrix \bar{V}_Ω and its complement \bar{V}_{Ω^c} . As we mention in Chapter 5, PASS Encrypt is inspired by a signature scheme, called PASS Sign, which also uses the partial Vandermonde matrix \bar{V}_Ω . It was originally introduced by Hoffstein et al. [HPS⁺14] and equipped with a rigorous security proof by Lu et al. [LZA18]. From a high level perspective, PASS Sign is an adaptation of Lyubashevsky’s Fiat-Shamir with aborts signature [Lyu09, Lyu12] to the partial Vandermonde setting. In contrast to PASS Encrypt, the signature scheme does not make use of the complement partial Vandermonde matrix \bar{V}_{Ω^c} and hence the novel attack that we

¹A deterministic PKE scheme cannot be IND-CPA secure as an adversary can simply encrypt both messages using the public key and decide which one is used in the challenge ciphertext.

²We briefly introduce the digital signature scheme PASS Sign in the introduction to Chapter 5. It uses mathematical concepts which are very similar to the ones of PASS Encrypt.

³In Fourier SIS, the partial Fourier matrix, instead of the partial Vandermonde matrix, is used.

present in this chapter does not apply to it. Nevertheless, it still benefits from the homomorphic properties of $\bar{\mathbf{V}}_\Omega$ which are used in a more recent work by Doröz et al. [DHSS20] to build a compact lattice-based aggregate signature scheme.

7.1.3 Roadmap

The rest of the chapter is structured as follows. We start in Section 7.2 with presenting the slightly modified version of PASS Encrypt, then prove its correctness and its security based on the hardness of decision PV-Knap and PASS. Furthermore, we demonstrate its additive and multiplicative homomorphic properties. Later, in Section 7.3 we provide an analysis of its security against known attacks before giving concrete sample parameters in Section 7.4 and a comparison with two other lattice-based encryption schemes.

7.2 PASS Encrypt

In the following, we use the same notation as in Chapter 5, that we quickly recall here. For simplicity, we focus on power-of-2 cyclotomics. Let ν be a power of 2 and q be a prime such that $q \equiv 1 \pmod{\nu}$. Further, we set $n = \nu/2$ and $t = n/2$. There are exactly n primitive ν -th roots of unity over \mathbb{Z}_q and for the key generation of our scheme we need to choose at random t (i.e., half) of them. To this end, we denote by \mathcal{P}_t the set of all subsets Ω of size t of all primitive ν -th roots of unity over \mathbb{Z}_q . Every Ω defines the corresponding partial Vandermonde matrix $\bar{\mathbf{V}}_\Omega$ and its complement $\bar{\mathbf{V}}_{\Omega^c}$. Let ψ_f, ψ_r, ψ_s be distributions over \mathbb{Z}^n . Recall that we denote by $+$ and \circ component-wise addition and multiplication of vectors over \mathbb{Z}_q . The message space \mathbf{M} is given by $\{0, 1\}^n$, and we select a message $\mathbf{m} \in \mathbf{M}$. Finally, let p be a small prime which is coprime to q and let λ denote the security parameter.

Protocol 7.1: Our slightly modified version of PASS Encrypt.

KGen(1^λ).	Sample $\mathbf{f} \leftarrow \psi_f$ and $\Omega \leftarrow U(\mathcal{P}_t)$, return $\text{sk} = \mathbf{f} \in \mathbb{Z}^n$, $\text{pk} = (\Omega, \bar{\mathbf{V}}_\Omega \mathbf{f}) \in \mathbb{Z}_q^t \times \mathbb{Z}_q^t$.
Enc(pk, \mathbf{m}).	Sample $\mathbf{r} \leftarrow \psi_r, \mathbf{s} \leftarrow \psi_s$, set $\mathbf{r}' = \mathbf{p}\mathbf{r}$ and $\mathbf{m}' = \mathbf{p}\mathbf{s} + \mathbf{m}$, set $\mathbf{e} = (\bar{\mathbf{V}}_\Omega \mathbf{r}' \circ \text{pk}) + \bar{\mathbf{V}}_\Omega \mathbf{m}'$ set $\mathbf{e}' = \bar{\mathbf{V}}_{\Omega^c} \mathbf{r}'$, set $\mathbf{e}'' = \bar{\mathbf{V}}_{\Omega^c} \mathbf{m}'$, return $\mathbf{c} = (\mathbf{e}, \mathbf{e}', \mathbf{e}'') \in \mathbb{Z}_q^t \times \mathbb{Z}_q^{n-t} \times \mathbb{Z}_q^{n-t}$.
Dec(sk, \mathbf{c}).	Compute $\mathbf{c}' = (\mathbf{e}' \circ \bar{\mathbf{V}}_{\Omega^c} \text{sk}) + \mathbf{e}'' \in \mathbb{Z}_q^{n-t}$, using knowledge of Ω and Ω^c combine \mathbf{e} and \mathbf{c}' to obtain \mathbf{c}'' as a vector over \mathbb{Z}_q^n , return $\bar{\mathbf{V}}^{-1} \mathbf{c}'' \pmod{p}$.

We now describe our slightly modified version of PASS Encrypt, as summarized in Protocol 7.1. During key generation, we sample \mathbf{f} from the distribution ψ_f , defining the secret key $\text{sk} = \mathbf{f} \in \mathbb{Z}^n$. Then, we sample $\Omega \in \mathbb{Z}_q^t$ uniformly at random over the set \mathcal{P}_t , which determines the public key $\text{pk} = (\Omega, \bar{\mathbf{V}}_\Omega \mathbf{f}) \in \mathbb{Z}_q^t \times \mathbb{Z}_q^t$, where the second part is the partial Vandermonde transform of \mathbf{f} evaluated at the roots given by Ω . In order to encrypt a message $\mathbf{m} \in \mathbf{M}$, we sample two random vectors $\mathbf{r} \leftarrow \psi_r$ and $\mathbf{s} \leftarrow \psi_s$, which define $\mathbf{r}' = \mathbf{p}\mathbf{r}$ and $\mathbf{m}' = \mathbf{p}\mathbf{s} + \mathbf{m}$. This randomizes the

message vector and thus converts PASS Encrypt from a deterministic in a randomized scheme. The ciphertext \mathbf{c} is then given by three elements. The first is $\mathbf{e} = (\overline{\mathbf{V}}_{\Omega} \mathbf{r}' \circ \text{pk}) + \overline{\mathbf{V}}_{\Omega} \mathbf{m}' \in \mathbb{Z}_q^t$, using the partial Vandermonde matrix $\overline{\mathbf{V}}_{\Omega}$. And the other two are given by $\mathbf{e}' = \overline{\mathbf{V}}_{\Omega^c} \mathbf{r}' \in \mathbb{Z}_q^{n-t}$ and $\mathbf{e}'' = \overline{\mathbf{V}}_{\Omega^c} \mathbf{m}' \in \mathbb{Z}_q^{n-t}$, using the complementary partial Vandermonde matrix $\overline{\mathbf{V}}_{\Omega^c}$. In order to decrypt a ciphertext \mathbf{c} , we use the secret key sk to first compute $\mathbf{c}' = (\mathbf{e}' \circ \overline{\mathbf{V}}_{\Omega^c} \text{sk}) + \mathbf{e}'' \in \mathbb{Z}_q^{n-t}$. Now, using the knowledge of Ω and Ω^c , we can combine \mathbf{e} and \mathbf{c}' to obtain a full vector \mathbf{c}'' over \mathbb{Z}_q^n . The decryption algorithm then returns $\overline{\mathbf{V}}^{-1} \mathbf{c}'' \bmod p$. We give some sample parameters for PASS Encrypt in Section 7.4 in Table 7.1.

Our version of PASS Encrypt differs in two aspects from the original version as presented in [HS15, Sec. 4]. First, they use the partial Fourier transform (that they denote by \mathcal{F}_S) instead of the partial Vandermonde transform $\overline{\mathbf{V}}_{\Omega}$, and second, in their case it always yields $\mathbf{s} = \mathbf{0}$ and thus $\mathbf{m}' = \mathbf{m}$. This makes the third part \mathbf{e}'' of their ciphertext only dependent on \mathbf{m} and hence the scheme deterministic. Additionally, the modifications also apply to the second version of the original proposed scheme, see [HS15, Sec. 6].

7.2.1 Correctness

We now show that the PKE scheme defined above is perfectly correct under a proper choice of parameters. See Definition 1.19 for the formal statement of the correctness property of a PKE scheme.

Lemma 7.1 (Correctness)

Let \mathcal{P}_t, p and ψ_f, ψ_r, ψ_s be the public parameters of PASS Encrypt. Assume that there exist $\alpha, \beta > 0$ such that for $\mathbf{f} \leftarrow \psi_f, \mathbf{r} \leftarrow \psi_r$ and $\mathbf{s} \leftarrow \psi_s$ it yields with probability 1 that $\|\mathbf{f}\|_{\infty} \leq 1, \|\mathbf{r}\|_1 \leq \alpha$ and $\|\mathbf{s}\|_{\infty} \leq \beta$. Further, we require $p(\alpha + \beta) + 1 < q/2$. For every key pair $(\text{sk}, \text{pk}) \leftarrow \text{KGen}(1^\lambda)$ and message $\mathbf{m} \in M$, it holds

$$\Pr [\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, \mathbf{m})) = \mathbf{m}] = 1.$$

Proof: The decryption oracle first computes $\mathbf{c}' = (\mathbf{e}' \circ \overline{\mathbf{V}}_{\Omega^c} \text{sk}) + \mathbf{e}'' \in \mathbb{Z}_q^{n-t}$ and then combines it with $\mathbf{e} \in \mathbb{Z}_q^t$ in order to obtain a full vector $\mathbf{c}'' \in \mathbb{Z}_q^n$. To guarantee correctness, we need to make sure that $\overline{\mathbf{V}}^{-1} \mathbf{c}'' = \mathbf{m} \bmod p$. Using the definition of sk, \mathbf{e}' and \mathbf{e}'' it yields

$$\mathbf{c}' = (\overline{\mathbf{V}}_{\Omega^c} \mathbf{r}' \circ \overline{\mathbf{V}}_{\Omega^c} \mathbf{f}) + \overline{\mathbf{V}}_{\Omega^c} \mathbf{m}' = \overline{\mathbf{V}}_{\Omega^c} (\mathbf{r}' \cdot \mathbf{f} + \mathbf{m}').$$

Simultaneously, using the definition of pk and \mathbf{e} it yields

$$\mathbf{e} = (\overline{\mathbf{V}}_{\Omega} \mathbf{r}' \circ \overline{\mathbf{V}}_{\Omega} \mathbf{f}) + \overline{\mathbf{V}}_{\Omega} \mathbf{m}' = \overline{\mathbf{V}}_{\Omega} (\mathbf{r}' \cdot \mathbf{f} + \mathbf{m}').$$

In both cases we use the multiplicative and homomorphic properties of $\overline{\mathbf{V}}_{\Omega}$ and $\overline{\mathbf{V}}_{\Omega^c}$ as presented in Section 5.2.1. Thus, combining both \mathbf{c}' and \mathbf{e} provides $\mathbf{c}'' = \overline{\mathbf{V}} (\mathbf{r}' \cdot \mathbf{f} + \mathbf{m}')$ and thus $\overline{\mathbf{V}}^{-1} \mathbf{c}'' = \mathbf{r}' \cdot \mathbf{f} + \mathbf{m}' = \text{pr} \cdot \mathbf{f} + \text{ps} + \mathbf{m} \bmod q$. Hence, $\overline{\mathbf{V}}^{-1} \mathbf{c}'' \bmod p = \mathbf{m} \bmod p$ if $\|\text{pr} \cdot \mathbf{f} + \text{ps} + \mathbf{m}\|_{\infty} < q/2$. Using the properties of power-of-2 cyclotomics to bound the infinity norm of the product of two elements as presented in Example 1.2 and that $\mathbf{m} \in M = \{0, 1\}^n$,

it yields

$$\begin{aligned} \|\mathbf{pr} \cdot \mathbf{f} + \mathbf{ps} + \mathbf{m}\|_\infty &\leq p\|\mathbf{r} \cdot \mathbf{f}\|_\infty + p\|\mathbf{s}\|_\infty + \|\mathbf{m}\|_\infty \\ &\leq p \cdot \|\mathbf{f}\|_\infty \cdot \|\mathbf{r}\|_1 + p\|\mathbf{s}\|_\infty + \|\mathbf{m}\|_\infty \\ &\leq p\alpha + p\beta + 1. \end{aligned}$$

As we require $p(\alpha+\beta)+1 < q/2$, the decryption algorithm decrypts correctly with probability 1. ■

7.2.2 Security

In this section, we prove the security of PASS Encrypt as defined above based on the hardness of PASS and dec-PV-Knap, both problems are defined in Section 5.2. We use the standard notion of IND-CPA security whose proper definition we recall in Definition 1.20.

In order to show the IND-CPA security of PASS Encrypt, we use a common game-hopping argument, as summarized in Protocol 7.2. Game 1 corresponds to the proposed PASS Encrypt. In Game 2 we change the definition of \mathbf{e} , in Game 3 the one of \mathbf{e}' and last in Game 4 the one of \mathbf{e}'' . Note that in Game 2, 3 and 4 the decryption algorithm does in general not succeed as the ciphertext parts \mathbf{e}, \mathbf{e}' or/and \mathbf{e}'' , when chosen uniformly at random, do in general not possess a small preimage under $\overline{\mathbf{V}}_\Omega$ or $\overline{\mathbf{V}}_{\Omega^c}$, respectively. For the proof of IND-CPA security, however, this does not pose any problem.

Protocol 7.2: Game hopping for IND-CPA security of PASS Encrypt.

	Game 1	Game 2
KGen:	$\text{sk} = \mathbf{f} \leftarrow \psi_f, \Omega \leftarrow U(\mathcal{P}_t)$ $\text{pk} = (\Omega, \overline{\mathbf{V}}_\Omega \mathbf{f} \bmod q)$	$\text{sk} = \mathbf{f} \leftarrow \psi_f, \Omega \leftarrow U(\mathcal{P}_t)$ $\text{pk} = (\Omega, \overline{\mathbf{V}}_\Omega \mathbf{f} \bmod q)$
Enc:	$\mathbf{r} \leftarrow \psi_r, \mathbf{s} \leftarrow \psi_s$ $\mathbf{r}' = \mathbf{pr}, \mathbf{m}' = \mathbf{ps} + \mathbf{m}$ $\mathbf{e} = (\overline{\mathbf{V}}_\Omega \mathbf{r}' \circ \text{pk}) + \overline{\mathbf{V}}_\Omega \mathbf{m}'$ $\mathbf{e}' = \overline{\mathbf{V}}_{\Omega^c} \mathbf{r}'$ $\mathbf{e}'' = \overline{\mathbf{V}}_{\Omega^c} \mathbf{m}'$	$\mathbf{r} \leftarrow \psi_r, \mathbf{s} \leftarrow \psi_s$ $\mathbf{r}' = \mathbf{pr}, \mathbf{m}' = \mathbf{ps} + \mathbf{m}$ $\mathbf{e} \leftarrow U(\mathbb{Z}_q^t)$ $\mathbf{e}' = \overline{\mathbf{V}}_{\Omega^c} \mathbf{r}'$ $\mathbf{e}'' = \overline{\mathbf{V}}_{\Omega^c} \mathbf{m}'$
	Game 3	Game 4
KGen:	$\text{sk} = \mathbf{f} \leftarrow \psi_f, \Omega \leftarrow U(\mathcal{P}_t)$ $\text{pk} = (\Omega, \overline{\mathbf{V}}_\Omega \mathbf{f} \bmod q)$	$\text{sk} = \mathbf{f} \leftarrow \psi_f, \Omega \leftarrow U(\mathcal{P}_t)$ $\text{pk} = (\Omega, \overline{\mathbf{V}}_\Omega \mathbf{f} \bmod q)$
Enc:	$\mathbf{r} \leftarrow \psi_r, \mathbf{s} \leftarrow \psi_s$ $\mathbf{r}' = \mathbf{pr}, \mathbf{m}' = \mathbf{ps} + \mathbf{m}$ $\mathbf{e} \leftarrow U(\mathbb{Z}_q^t)$ $\mathbf{e}' \leftarrow U(\mathbb{Z}_q^{n-t})$ $\mathbf{e}'' = \overline{\mathbf{V}}_{\Omega^c} \mathbf{m}'$	$\mathbf{r} \leftarrow \psi_r, \mathbf{s} \leftarrow \psi_s$ $\mathbf{r}' = \mathbf{pr}, \mathbf{m}' = \mathbf{ps} + \mathbf{m}$ $\mathbf{e} \leftarrow U(\mathbb{Z}_q^t)$ $\mathbf{e}' \leftarrow U(\mathbb{Z}_q^{n-t})$ $\mathbf{e}'' \leftarrow U(\mathbb{Z}_q^{n-t})$

Lemma 7.2 (Security)

Let \mathcal{P}_t, p and ψ_f, ψ_r, ψ_s be the public parameters of PASS Encrypt and the message $\mathbf{m} \in \mathcal{M}$. Assuming the hardness $\text{dec-PV-Knap}_{\psi_1}$, $\text{dec-PV-Knap}_{\psi_2}$ and $\text{PASS}_{\psi_f, \psi_1, \psi_2}$, where $\psi_1 = p \cdot \psi_r$ and $\psi_2 = p \cdot \psi_s + \mathbf{m}$, the encryption scheme as summarized in Protocol 7.1 is IND-CPA secure.

Proof: Note that Game 1 corresponds to the proposed PASS Encrypt. The only difference between Game 1 and Game 2 is the way how \mathbf{e} is defined. In the first game, it is a partial Vandermonde transform, given by $(\overline{\mathbf{V}}_{\Omega}\mathbf{r}' \circ \text{pk}) + \overline{\mathbf{V}}_{\Omega}\mathbf{m}' = \overline{\mathbf{V}}_{\Omega}(\mathbf{r}' \cdot \mathbf{f} + \mathbf{m}')$, and in the second game it is sampled uniformly at random over \mathbb{Z}_q^t . Notice that pk, \mathbf{e}' and \mathbf{e}'' are *not* independent from \mathbf{e} , but assuming the hardness of $\text{PASS}_{\psi_f, \psi_1, \psi_2}$, with $\psi_1 = p \cdot \psi_r$ and $\psi_2 = p \cdot \psi_s + \mathbf{m}$, an adversary cannot distinguish between the two games.

Now, we are studying the difference between Game 2 and Game 3. Here, the second ciphertext part \mathbf{e}' is replaced by a uniform element over \mathbb{Z}_q^{n-t} . We remark, that \mathbf{e}' is independent from the other two ciphertext parts \mathbf{e} and \mathbf{e}'' and also independent from the secret key. Thus, assuming the hardness of $\text{dec-PV-Knap}_{\psi_1}$, an adversary cannot distinguish Game 2 from Game 3.

The only difference between Game 3 and Game 4 is the definition of \mathbf{e}'' . With the same argument, they cannot distinguish Game 3 from Game 4, assuming the hardness of $\text{dec-PV-Knap}_{\psi_2}$.

In the last Game 4, the ciphertext $\mathbf{c} = (\mathbf{e}, \mathbf{e}', \mathbf{e}'')$ is independent of the message \mathbf{m} and the key pair (sk, pk) . Thus, the adversary has no chance to distinguish two ciphertexts in the IND-CPA security game better than to guess. ■

We would like to emphasize the following connection of PASS Encrypt to ideal lattices. As elaborated in Section 5.2.1, the public key of PASS Encrypt given by the partial Vandermonde transform $\text{pk} = \overline{\mathbf{V}}_{\Omega}\mathbf{f}$ can be seen as a way to specify the coset $\mathbf{f} + \mathcal{I}$, where $\mathcal{I} = \prod_{\omega_j \in \Omega} \langle q, x - \omega_j \rangle$ is an ideal lattice. Simultaneously, the complement partial Vandermonde transforms $\overline{\mathbf{V}}_{\Omega_c}\mathbf{r}'$ and $\overline{\mathbf{V}}_{\Omega_c}\mathbf{m}'$ (i.e., the second and third part of the ciphertext of PASS Encrypt) can be seen as a way to specify the cosets $\mathbf{r}' + \mathcal{I}'$ and $\mathbf{m}' + \mathcal{I}'$, where $\mathcal{I} \cdot \mathcal{I}' = \langle q \rangle$. In other words, PASS Encrypt allows a formulation directly in the language of ideal lattices.

7.2.3 Homomorphic Properties

In the following, we show that our slight modifications on PASS Encrypt preserve its additive and multiplicative homomorphic properties, as originally demonstrated by Hoffstein and Silverman [HS15, Sec. 5].

Additive Homomorphic

For addition, we can simply sum the different parts of two given ciphertexts to obtain the encryption of the sum of the original messages. To decrypt the sum, we can use the same decryption algorithm as for a single ciphertext.

More precisely, given for a fixed key pair (sk, pk) two ciphertexts $\mathbf{c}_1 = (\mathbf{e}_1, \mathbf{e}'_1, \mathbf{e}''_1)$ and $\mathbf{c}_2 = (\mathbf{e}_2, \mathbf{e}'_2, \mathbf{e}''_2)$ on two messages \mathbf{m}_1 and \mathbf{m}_2 , where during encryption the random ring elements $\mathbf{r}_1, \mathbf{s}_1$ and $\mathbf{r}_2, \mathbf{s}_2$ were used. Then, the element $\mathbf{c} = (\mathbf{e}_1 + \mathbf{e}_2, \mathbf{e}'_1 + \mathbf{e}'_2, \mathbf{e}''_1 + \mathbf{e}''_2)$ defines the ciphertext of the message $\mathbf{m} = \mathbf{m}_1 + \mathbf{m}_2$ with encryption randomness $\mathbf{r} = \mathbf{r}_1 + \mathbf{r}_2$ and $\mathbf{s} = \mathbf{s}_1 + \mathbf{s}_2$. Here, we only use the linearity of matrix-vector products.

Multiplicative Homomorphic

The situation is slightly more complex for multiplication, where an additional cross term has to be provided in the ciphertext in order to enable the decryption of the product of two ciphertexts.

In more details, assume that we are given for a fixed key pair (sk, pk) two ciphertexts $\mathbf{c}_1 = (\mathbf{e}_1, \mathbf{e}'_1, \mathbf{e}''_1)$ and $\mathbf{c}_2 = (\mathbf{e}_2, \mathbf{e}'_2, \mathbf{e}''_2)$ on two messages \mathbf{m}_1 and \mathbf{m}_2 , where during encryption the random

elements $\mathbf{r}_1, \mathbf{s}_1$ and $\mathbf{r}_2, \mathbf{s}_2$ were used. In order to provide enough information to recover the product message $\mathbf{m}_1 \cdot \mathbf{m}_2 \bmod p$, we need to transmit in the ciphertext the respective products $\mathbf{e} = \mathbf{e}_1 \circ \mathbf{e}_2$, $\mathbf{e}' = \mathbf{e}'_1 \circ \mathbf{e}'_2$ and $\mathbf{e}'' = \mathbf{e}''_1 \circ \mathbf{e}''_2$, and additionally a cross term $\mathbf{E} = \mathbf{e}'_1 \circ \mathbf{e}''_2 + \mathbf{e}'_2 \circ \mathbf{e}''_1$. To decrypt, we use $\mathbf{e}, \mathbf{e}', \mathbf{e}'', \mathbf{E}$ and the secret key \mathbf{sk} to compute

$$\begin{aligned} \mathbf{c}' &= (\mathbf{e}' \circ (\overline{\mathbf{V}}_{\Omega^c} \mathbf{sk})^2) + (\mathbf{E} \circ \overline{\mathbf{V}}_{\Omega^c} \mathbf{sk}) + \mathbf{e}'' \\ &= (\overline{\mathbf{V}}_{\Omega^c} \mathbf{r}'_1 \circ \overline{\mathbf{V}}_{\Omega^c} \mathbf{r}'_2 \circ (\overline{\mathbf{V}}_{\Omega^c} \mathbf{f})^2) + (\overline{\mathbf{V}}_{\Omega^c} \mathbf{r}'_1 \circ \overline{\mathbf{V}}_{\Omega^c} \mathbf{m}'_2 \circ \overline{\mathbf{V}}_{\Omega^c} \mathbf{f}) \\ &\quad + (\overline{\mathbf{V}}_{\Omega^c} \mathbf{r}'_2 \circ \overline{\mathbf{V}}_{\Omega^c} \mathbf{m}'_1 \circ \overline{\mathbf{V}}_{\Omega^c} \mathbf{f}) + (\overline{\mathbf{V}}_{\Omega^c} \mathbf{m}'_1 \circ \overline{\mathbf{V}}_{\Omega^c} \mathbf{m}'_2) \\ &= \overline{\mathbf{V}}_{\Omega^c} ((\mathbf{r}'_1 \mathbf{r}'_2 \mathbf{f}^2) + (\mathbf{r}'_1 \mathbf{m}'_2 \mathbf{f}) + (\mathbf{r}'_2 \mathbf{m}'_1 \mathbf{f}) + (\mathbf{m}'_1 \mathbf{m}'_2)). \end{aligned}$$

On the other hand, it yields

$$\begin{aligned} \mathbf{e} &= \mathbf{e}_1 \circ \mathbf{e}_2 \\ &= ((\overline{\mathbf{V}}_{\Omega} \mathbf{r}'_1 \circ \mathbf{pk}) + \overline{\mathbf{V}}_{\Omega} \mathbf{m}'_1) \circ ((\overline{\mathbf{V}}_{\Omega} \mathbf{r}'_2 \circ \mathbf{pk}) + \overline{\mathbf{V}}_{\Omega} \mathbf{m}'_2) \\ &= (\overline{\mathbf{V}}_{\Omega} (\mathbf{r}'_1 \mathbf{f} + \mathbf{m}'_1)) \circ (\overline{\mathbf{V}}_{\Omega} (\mathbf{r}'_2 \mathbf{f} + \mathbf{m}'_2)) \\ &= \overline{\mathbf{V}}_{\Omega} ((\mathbf{r}'_1 \mathbf{r}'_2 \mathbf{f}^2) + (\mathbf{r}'_1 \mathbf{m}'_2 \mathbf{f}) + (\mathbf{r}'_2 \mathbf{m}'_1 \mathbf{f}) + (\mathbf{m}'_1 \mathbf{m}'_2)). \end{aligned}$$

Combining \mathbf{c}' and \mathbf{e} gives the full Vandermonde transform and by applying $\overline{\mathbf{V}}^{-1}$, we obtain $\mathbf{c}'' = (\mathbf{r}'_1 \mathbf{r}'_2 \mathbf{f}^2) + (\mathbf{r}'_1 \mathbf{m}'_2 \mathbf{f}) + (\mathbf{r}'_2 \mathbf{m}'_1 \mathbf{f}) + (\mathbf{m}'_1 \mathbf{m}'_2) \bmod q$. If $\|\mathbf{c}''\|_{\infty} < q/2$, we can compute $\mathbf{c}'' \bmod p = \mathbf{m}'_1 \mathbf{m}'_2 \bmod p = \mathbf{m}_1 \mathbf{m}_2 \bmod p$. Here, we use that $\mathbf{r}'_j = \mathbf{p} \mathbf{r}_j$ and $\mathbf{m}'_j = \mathbf{p} \mathbf{s}_j + \mathbf{m}_j$ for $j \in \{1, 2\}$.

7.3 Security Analysis

We start this section by investigating the concrete security of PASS Encrypt against best known attacks. We present three attacks that we call the Key Recovery (Section 7.3.1), the Randomness Recovery (Section 7.3.2) and the Plaintext Recovery Using Hints Attacks (Section 7.3.3). The first two attacks are already considered in the original PASS Encrypt proposal by Hoffstein and Silverman [HS15]. We restate them for completeness and rephrase them in the primal attack framework of LWE as done by Alkim et al. [ADPS16]. Further, [HS15] use the less common notion of MIPS-years, where one MIPS-year equals the number of instructions executed during one year of computing at one million instructions per second. Note that in the parallel line of work concerning PASS Sign, the same type of attacks is studied as well ([HPS⁺14, LZA18, DHSS20]).

Essentially, recovering the secret key \mathbf{sk} (resp. the randomness \mathbf{r}') used in the encryption algorithm (see Protocol 7.1) corresponds to solve an instance of PV-Knap with regard to the partial Vandermonde matrix $\overline{\mathbf{V}}_{\Omega}$ (resp. the complement partial Vandermonde matrix $\overline{\mathbf{V}}_{\Omega^c}$). However, no attack that aims at recovering the secret vector of the PASS instance given by a ciphertext has been studied so far. This leads us to the third attack, that we call the Plaintext Recovery Using Hints Attack in the following. This novel attack takes the design of PASS Encrypt into account and thus improves our understanding of its security. We then argue why we move from the Fourier to the Vandermonde transform in Section 7.3.4. Furthermore, we show how to interpret the partial Vandermonde LWE and SIS problems in terms of error-correcting codes in Section 7.3.5.

Finally, we give concrete sample parameters and security estimates in Section 7.4 and compare PASS Encrypt with two other efficient lattice-based PKE schemes in Section 7.4.2.

7.3.1 Key Recovery Attack

We now describe the first attack against PASS Encrypt, as already considered in the original proposal [HS15, Sec. 7]. We restate it for completeness and rephrase it in the attack framework of LWE as done by Alkim et al. [ADPS16], using the BKZ algorithm with quantum sieving to solve the associated Unique Shortest Vector Problem (u-SVP, Def. 1.3). The second component of the public key pk of PASS Encrypt is a vector $\mathbf{g} \in \mathbb{Z}_q^t$ defined as $\mathbf{g} = \overline{\mathbf{V}}_\Omega \mathbf{f} \bmod q$, where $\mathbf{f} \leftarrow \psi_f$. We can write $\overline{\mathbf{V}}_\Omega = [\mathbf{A}|\mathbf{B}] \in \mathbb{Z}_q^{t \times n}$ with $\mathbf{A} \in \mathbb{Z}_q^{t \times (n-t)}$ and $\mathbf{B} \in \mathbb{Z}_q^{t \times t}$, where \mathbf{B} has full rank t and thus by multiplying $\overline{\mathbf{V}}_\Omega$ by the inverse of \mathbf{B} , it takes the form $[\tilde{\mathbf{A}}|\mathbf{I}_t] \in \mathbb{Z}_q^{t \times n}$, for some matrix $\tilde{\mathbf{A}} \in \mathbb{Z}_q^{t \times (n-t)}$ and \mathbf{I}_t the identity matrix of order t . Hence, we can transform the equation above to

$$\tilde{\mathbf{g}} = \tilde{\mathbf{A}} \cdot \mathbf{f}_1 + \mathbf{f}_2 \bmod q, \quad (7.1)$$

where $\tilde{\mathbf{g}} = \mathbf{B}^{-1} \mathbf{g}$ and $\mathbf{f} = (\mathbf{f}_1, \mathbf{f}_2)^T$ with $\mathbf{f}_1 \in \mathbb{Z}^{n-t}$ and $\mathbf{f}_2 \in \mathbb{Z}^t$. Equation 7.1 can be seen as an instance of LWE in its Hermite Normal Form (HNF), as defined in Section 1.4.1, with public matrix $\tilde{\mathbf{A}}$ of LWE dimension $n-t$ and with t denoting the number of given samples. In doing so, we ignore the known structure of the matrix $\tilde{\mathbf{A}}$ and treat it as a uniform random matrix. This is a common approach used in structured lattice-based cryptography as no cryptanalytic technique making use of the algebraic structure is known, for a more elaborated discussion see [ACD⁺18]. We proceed as we usually do for LWE (see [ADPS16] for more details) and rewrite the equation above as $\tilde{\mathbf{A}} \cdot \mathbf{f}_1 + \mathbf{f}_2 - \tilde{\mathbf{g}} = \mathbf{0} \bmod q$. This defines an instance of u-SVP in the lattice $\Lambda = \Lambda(\tilde{\mathbf{A}}, \tilde{\mathbf{g}})$ given by

$$\Lambda = \{(\mathbf{x}, \mathbf{y}, w)^T \in \mathbb{Z}^t \times \mathbb{Z}^{n-t} \times \mathbb{Z} : \mathbf{x} + \tilde{\mathbf{A}} \cdot \mathbf{y} - w\tilde{\mathbf{g}} = \mathbf{0} \bmod q\}.$$

A basis of this lattice is given by the column vectors of

$$\mathbf{C} = \begin{bmatrix} q\mathbf{I}_t & -\tilde{\mathbf{A}} & \tilde{\mathbf{g}} \\ \mathbf{0}_{(n-t) \times t} & \mathbf{I}_{n-t} & \mathbf{0}_{(n-t) \times 1} \\ \mathbf{0}_{1 \times t} & \mathbf{0}_{1 \times (n-t)} & 1 \end{bmatrix} \in \mathbb{Z}^{(n+1) \times (n+1)},$$

where for $n, m \in \mathbb{N}$, we denote by $\mathbf{0}_{n \times m}$ the $n \times m$ matrix composed of zeros. The lattice Λ has full rank $n+1$ as it has an upper triangular form. Further, its determinant is q^t , corresponding to the determinant of \mathbf{C} . It is easy to see that the vector $(\mathbf{f}_2, \mathbf{f}_1, 1)^T \in \mathbb{Z}^{n+1}$ lies in Λ , where its norm depends on the distribution ψ_f . Assuming that $\psi_f = U(\{-1, 0, 1\})$, we expect its Euclidean norm to be $\approx \sqrt{\frac{2n}{3}}$.

When estimating the expected length of a shortest vector, one can use the Gaussian heuristic, as explained in Section 1.2. More precisely, the Gaussian heuristic for the given lattice Λ with determinant q^t and of dimension $n+1$ (see Equation 1.2) states that the shortest vector in Λ has Euclidean norm approximately $\sqrt{\frac{n+1}{2\pi e}} \cdot q^{\frac{t}{n+1}} \approx \sqrt{q \cdot n}$, for $t = n/2$. This is much larger than the norm of $(\mathbf{f}_2, \mathbf{f}_1, 1)^T$. In other words, we assume that $(\mathbf{f}_2, \mathbf{f}_1, 1)^T$ is the unique shortest vector and the second shortest vector has the norm following the Gaussian heuristic. This is an instance of u-SVP, as presented in Definition 1.3. If we assume that q is linear in n , then the ratio of the shortest and the second shortest vector (also called the SVP gap) is approximately $\frac{1}{\sqrt{n}}$.

The u-SVP instance can be solved by the BKZ algorithm and its running time depends on the used blocksize within BKZ. We use the publicly accessible⁴ Leaky LWE Estimator [DDGR20] to estimate the necessary blocksize for the BKZ algorithm, denoted as *bikz*. The algorithm BKZ itself uses an SVP oracle in dimension *bikz*. As in [ADPS16], we evaluate the running time of BKZ using

⁴<https://github.com/lducas/leaky-LWE-Estimator>

the core SVP hardness, thus considering only the cost of one call to an SVP oracle in dimension $bikz$.

As the best known heuristic quantum algorithm to solve SVP in dimension $bikz$ [Laa15] runs in time $2^{0.265 \cdot bikz}$, we give the number of quantum security bits by $0.265 \cdot bikz$. We give concrete sample parameters and values for this attack against PASS Encrypt in Section 7.4 in Table 7.1.

7.3.2 Randomness Recovery Attack

We now describe the second attack against PASS Encrypt, which aims at recovering the underlying randomness \mathbf{r} used during encryption. The second component of the ciphertext \mathbf{c} of PASS Encrypt is given by a vector $\mathbf{e}' \in \mathbb{Z}_q^{n-t}$ satisfying $\mathbf{e}' = \bar{\mathbf{V}}_{\Omega} \mathbf{r}' \bmod q$, with $\mathbf{r}' = \mathbf{p}\mathbf{r}$ for $\mathbf{r} \leftarrow \psi_r$. As p is publicly known and coprime to q , we can divide the above by p to obtain $\mathbf{g} := \mathbf{e}'/p = \bar{\mathbf{V}}_{\Omega} \mathbf{r} \bmod q$. As for the Key Recovery Attack from above, we can interpret this as an instance of LWE. More precisely, we write $\bar{\mathbf{V}}_{\Omega} \mathbf{c} = [\mathbf{C}|\mathbf{D}] \in \mathbb{Z}_q^{(n-t) \times n}$ with $\mathbf{C} \in \mathbb{Z}_q^{(n-t) \times t}$ and $\mathbf{D} \in \mathbb{Z}_q^{(n-t) \times (n-t)}$, where \mathbf{D} has full rank $n-t$. We multiply the equation above by the inverse of \mathbf{D} to obtain

$$\tilde{\mathbf{g}} = \tilde{\mathbf{C}} \cdot \mathbf{r}_1 + \mathbf{r}_2 \bmod q, \quad (7.2)$$

where $[\tilde{\mathbf{C}}|\mathbf{I}_{n-t}] = [\mathbf{C}|\mathbf{D}] \cdot \mathbf{D}^{-1}$, $\tilde{\mathbf{g}} = \mathbf{D}^{-1} \mathbf{g}$ and $\mathbf{r} = (\mathbf{r}_1, \mathbf{r}_2)^T$ with $\mathbf{r}_1 \in \mathbb{Z}^t$ and $\mathbf{r}_2 \in \mathbb{Z}^{n-t}$. In other words, Equation 7.2 describes an instance of LWE in Hermite Normal Form with public matrix $\tilde{\mathbf{C}}$ of dimension t , where $n-t$ samples are given. Note that the roles of the dimension and the number of samples are exactly the reversed roles as in the Key Recovery Attack. However, for $t = n/2$, the dimension and number of samples of the LWE instances are in both attacks the same. Further, for $\psi_r = \psi_f$, as we do in the sample parameters in Table 7.1, both attacks are equally hard.

In order to recover the plaintext \mathbf{m} from the ciphertext $\mathbf{c} = (\mathbf{e}, \mathbf{e}', \mathbf{e}'')$, an attacker can use the same approach as in the Key Recovery Attack to solve the associated LWE instance (obtaining \mathbf{r}) and to compute $\mathbf{r}' = \mathbf{p}\mathbf{r}$. They can then use \mathbf{r}' to compute $\bar{\mathbf{V}}_{\Omega} \mathbf{m}' = \mathbf{e} - (\bar{\mathbf{V}}_{\Omega} \mathbf{r}' \circ \mathbf{pk})$. By combining $\bar{\mathbf{V}}_{\Omega} \mathbf{m}'$ and $\mathbf{e}'' = \bar{\mathbf{V}}_{\Omega} \mathbf{m}'$ to the full discrete Vandermonde transform $\bar{\mathbf{V}} \mathbf{m}'$, one can multiply it by $\bar{\mathbf{V}}^{-1}$ to obtain $\mathbf{m}' = \mathbf{p}\mathbf{s} + \mathbf{m} \bmod q$. Finally, the plaintext message $\mathbf{m} \bmod p$ can be recovered by computing $\mathbf{m}' \bmod p$.

7.3.3 Plaintext Recovery Using Hints Attack

In the following we present a new attack against PASS Encrypt, which is inspired by the recent work of Dachman-Soled et al. [DDGR20] on exploiting hints that are given on the LWE secret or noise. The first part of the ciphertext \mathbf{c} of PASS Encrypt is given by $\mathbf{e} = (\bar{\mathbf{V}}_{\Omega} \mathbf{r}' \circ \mathbf{pk}) + \bar{\mathbf{V}}_{\Omega} \mathbf{m}'$, where $\mathbf{r}' = \mathbf{p}\mathbf{r}$, $\mathbf{pk} = \bar{\mathbf{V}}_{\Omega} \mathbf{f}$ and $\mathbf{m}' = \mathbf{p}\mathbf{s} + \mathbf{m}$ with $\mathbf{f} \leftarrow \psi_f$, $\mathbf{r} \leftarrow \psi_r$ and $\mathbf{s} \leftarrow \psi_s$. Using the homomorphic properties of $\bar{\mathbf{V}}_{\Omega}$ as presented in Section 5.2.1, we can rewrite \mathbf{e} as $\mathbf{e} = \bar{\mathbf{V}}_{\Omega} \cdot (\mathbf{f} \cdot \mathbf{r}' + \mathbf{m}')$. Recall that $\text{Rot}(\mathbf{f})$ denotes the matrix describing the multiplication by \mathbf{f} in the coefficient embedding, see Section 1.1.3. In matrix form this gives

$$\mathbf{e} = \bar{\mathbf{V}}_{\Omega} \cdot (\text{Rot}(\mathbf{f}) \cdot \mathbf{r}' + \mathbf{I}_n \cdot \mathbf{m}') = [\mathbf{A}|\bar{\mathbf{V}}_{\Omega}] \cdot \begin{pmatrix} \mathbf{r}' \\ \mathbf{m}' \end{pmatrix} \bmod q,$$

where $\mathbf{A} = \bar{\mathbf{V}}_{\Omega} \cdot \text{Rot}(\mathbf{f}) \in \mathbb{Z}_q^{t \times n}$. Note that we can compute \mathbf{A} by knowing the roots $\omega_{i_\ell} \in \Omega$ for $\ell \in [t]$ and the public key \mathbf{pk} , and not necessarily the secret key $\mathbf{sk} = \mathbf{f}$, as explained in Lemma 5.1. As $\bar{\mathbf{V}}_{\Omega}$ has full rank t , we can use Gauss elimination to transform this equation into

$$\tilde{\mathbf{e}} = [\mathbf{B}|\mathbf{I}_t] \cdot \begin{pmatrix} \mathbf{r}' \\ \mathbf{m}' \end{pmatrix} = \mathbf{B} \cdot \tilde{\mathbf{r}}' + \tilde{\mathbf{m}}' \bmod q,$$

with $\mathbf{B} \in \mathbb{Z}_q^{t \times (2n-t)}$, $\tilde{\mathbf{r}}'$ containing \mathbf{r}' and the first $n-t$ coefficients of \mathbf{m}' and $\tilde{\mathbf{m}}'$ containing the last t coefficients of \mathbf{m}' . This corresponds to an instance of LWE of dimension $2n-t$ and t the number of given samples, with \mathbf{B} the public matrix.

At first sight, one can see that the LWE instance is of much larger dimension than in the two previous attacks, and thus one may wonder why this attack should provide tighter security estimates. As we will see now, this is because of the additional information provided by the rest of the ciphertext. The second and third part of the ciphertext \mathbf{e}' and \mathbf{e}'' can be viewed as hints on \mathbf{r}' and \mathbf{m}' . To be more precise, $\mathbf{e}' = \overline{\mathbf{V}}_{\Omega^c} \cdot \mathbf{r}'$ and $\mathbf{e}'' = \overline{\mathbf{V}}_{\Omega^c} \cdot \mathbf{m}'$. This can be rewritten as

$$\mathbf{e}' = [\overline{\mathbf{V}}_{\Omega^c} | \mathbf{0}_{(n-t) \times n}] \cdot \begin{pmatrix} \mathbf{r}' \\ \mathbf{m}' \end{pmatrix} \bmod q, \quad \mathbf{e}'' = [\mathbf{0}_{(n-t) \times n} | \overline{\mathbf{V}}_{\Omega^c}] \cdot \begin{pmatrix} \mathbf{r}' \\ \mathbf{m}' \end{pmatrix} \bmod q.$$

Note that the vector $(\tilde{\mathbf{r}}', \tilde{\mathbf{m}}')^T$ is simply a re-labeling of the vector $(\mathbf{r}', \mathbf{m}')^T$. In the language of Dachman-Soled et al. [DDGR20] this corresponds to $2(n-t)$ modular hints.

For simplicity, we assume that $\mathbf{m} = \mathbf{0}$ and thus $\mathbf{m}' = ps$. As p is a public parameter we can assume that the secret \mathbf{r}' and the noise \mathbf{m}' of the corresponding LWE sample are drawn from the distributions ψ_r and ψ_s , respectively.

As in the Key Recovery Attack, the number of security bits claims that a quantum algorithm would need at least a running time of $2^{0.265 \cdot \text{bikz}}$, where bikz is the blocksize resulting from the Leaky LWE Estimator [DDGR20].

7.3.4 Choice of Ring

In the original description of PASS Encrypt [HS15], the partial Fourier transform is used, and not as we propose the partial Vandermonde transform. The main difference between the original and our version is that the latter works over the ring of integers of some cyclotomic number field, whereas the first one works over the cyclic ring $\mathbb{Z}[x]/\langle x^n - 1 \rangle$, for some prime n . The setting in [HS15] is the following. Let n and q be primes satisfying $q \equiv 1 \pmod n$ and let ω be a primitive n -th root of unity in \mathbb{Z}_q . Further, let S be a subset of $[n]$ of size t and let $\Omega = \{\omega^{k-1} : k \in S\}$ and $\Omega^c = \{\omega^{k-1} : k \in [n] \setminus S\}$. The (discrete) partial Fourier transformation matrix $\overline{\mathbf{F}}_{\Omega}$ is defined as $\overline{\mathbf{F}}_{\Omega} := (\omega_j^{k-1})_{j \in [t], k \in [n]}$, where $\omega_j \in \Omega$ for $j \in [t]$. In an analogue manner to Section 5.2.1, where we define Partial Vandermonde SIS (Definition 5.3), we can define Partial Fourier SIS, denoted by PF-SIS. More concretely, for a given parameter $\beta > 0$, PF-SIS $_{\beta}$ asks to find an element $\mathbf{a} \in \mathbb{Z}[x]/\langle x^n - 1 \rangle$ of norm $\|\mathbf{a}\|_2 \leq \beta$ satisfying $\overline{\mathbf{F}}_{\Omega} \cdot \mathbf{a} = \mathbf{0} \bmod q$.

Lemma 7.3 (Solution to Partial Fourier SIS)

The problem PF-SIS $_{\beta}$ is easy to solve for any $\beta \geq \sqrt{n}$.

Proof: Let $\mathbf{1}$ denote the element in $\mathbb{Z}[x]/\langle x^n - 1 \rangle$ whose coefficient vector is given by $(1, \dots, 1)^T \in \mathbb{Z}^n$. The polynomial $x^n - 1$ can be factorized in the product $x^n - 1 = (x-1)(x^{n-1} + \dots + x + 1)$. As for any $j \in [t]$ the element ω_j is a solution to the equation $x^n - 1 \bmod q$, we also know that $\sum_{k \in [n]} \omega_j^{k-1} = 0 \bmod q$ and thus $\overline{\mathbf{F}}_{\Omega} \mathbf{1} = \mathbf{0} \bmod q$ with $\|\mathbf{1}\|_2 = \sqrt{n} \leq \beta$. ■

This generic solution only holds for the *homogeneous* problem PF-SIS, and not for the *inhomogeneous* Knapsack counterpart. However, we prefer to move to the ring $R = \mathbb{Z}[x]/\langle x^n + 1 \rangle$, where n is a power of two. For this ring, the so-called evaluation-at-1 attack does not work. Note

that the evaluating-at-1 approach already led to successful attacks against NTRU.

7.3.5 Code-Based Attack

We now explain how the partial Vandermonde transform can be interpreted in terms of error-correcting codes. For a gentle introduction to coding theory, we refer to the book of Roth [Rot06].

More formally, the partial Vandermonde transformation matrix $\bar{\mathbf{V}}_\Omega \in \mathbb{Z}_q^{t \times n}$, as defined in Section 5.2.1, describes the check matrix of an $[n, t, d]$ Reed-Solomon code, with $d = n - t + 1$ its minimum distance. Using the duality connection between PV-Knap and PV-LWE from Section 5.3, the PV-LWE matrix $\bar{\mathbf{V}}_{\text{inv}(\Omega^c)}^T \in \mathbb{Z}_q^{n \times (n-t)}$ corresponds to the generating matrix of the same code. More concretely, it is a punctured Reed-Solomon code fulfilling the optimal singleton-bound. Thus, it is a maximum distance separable code with correction capability $\lfloor (d-1)/2 \rfloor = \lfloor (n-t)/2 \rfloor$. Our typical choice of t is $t = \lfloor n/2 \rfloor$, and thus the correction capability is bounded above by $\lfloor n/4 \rfloor$. It is further a unit-derived code and the theory of error-decoding pairs provides efficient decoder, see for instance the work by Hurley and Hurley [HH18].

In order to prevent coding-based attacks, we need to choose distributions for the secret in PV-Knap and for the noise in PV-LWE such that the expected Hamming weight is not near the correction capability bound. For example, setting the distribution ψ as the uniform distribution over $\{-1, 0, 1\}^n$ and sampling $\mathbf{e} \leftarrow \psi$, we expect the Hamming weight of \mathbf{e} to be $2n/3$, which is much larger than the error-correcting capacity, which is at most $\lfloor n/4 \rfloor$. So this would be a choice that does not allow for code-based attacks. If, however, we set the distribution ψ as a sparse distribution over $\{-1, 0, 1\}^n$, where an element $\mathbf{e} \leftarrow \psi$ has only very few non-zero coefficients, let's say about $n/4$ non-zero coefficients, then the efficient decoder from [HH18] would apply.

7.4 Concrete Parameters

We propose the following sample parameters for PASS Encrypt, under which the scheme as presented in Protocol 7.1 is correct (Lemma 7.1). We consider the case where K is the ν -th cyclotomic number field with ν a power of 2. By $n = \nu/2$ we denote its degree and the number of rows of the partial Vandermonde matrix t is given by $n/2$. In Section 7.4.1 below we argue why this is the optimal choice for t . The parameter q denotes the modulus over which the partial Vandermonde transformation matrix $\bar{\mathbf{V}}_\Omega$ is taken. We require $q = 1 \pmod{\nu}$ such that the defining polynomial of K , given by $x^n + 1$, fully splits modulo q . Concretely, we provide two parameter sets, as summarized in Table 7.1. In the first, we choose $\nu = 2048$ and $q = 12289$, and in the second, we keep the same q and set $\nu = 4096$. Note that the relevant parameter for security is t .

We set the distributions ψ_f, ψ_r and ψ_s as $U(T_n(d))$, the uniform distribution over ternary polynomials with exactly d coefficients that equal 1, and d coefficients that equal -1 , and $n - 2d$ coefficients that equal 0, where $d = \lfloor n/3 \rfloor$. Thus, for every element $\mathbf{f} \leftarrow \psi_f$ (resp. $\mathbf{r} \leftarrow \psi_r$ and $\mathbf{s} \leftarrow \psi_s$) it yields $\|\mathbf{f}\|_\infty \leq 1$ (resp. $\|\mathbf{r}\|_1 \leq 2n/3$ and $\|\mathbf{s}\|_\infty \leq 1$) with probability 1. Hence, we can set the parameter α to $2n/3$ and β to 1. Fixing the number of coefficients that equal -1 and 1 makes it possible to set $\alpha = 2n/3$ (in order to keep *perfect* correctness for the given q), but adds a structural hint, as exploited by Dachman-Soled et al. [DDGR20, Sec. 6.3]. This structural hint roughly decreases the estimated *bikz* by 1. Further, we set p as 2.

We then provide the needed block sizes of the BKZ algorithm in order to perform the three attacks on PASS Encrypt for both parameter sets, as presented in Section 7.3. All estimations are computed with SageMath using the Leaky LWE Estimator [DDGR20].

Table 7.1: Sample parameters and security estimations for PASS Encrypt. The number of quantum security bits is computed as $\lfloor 0.265 \cdot \min_{j \in [3]}(bikz_j) \rfloor$.

Parameter	Set 1	Set 2
ν	2048	4096
n	1024	2048
t	512	1024
q	12289	12289
p	2	2
α	$\lfloor 2n/3 \rfloor$	$\lfloor 2n/3 \rfloor$
β	1	1
$\psi_f = \psi_r = \psi_s$	$U(T_n(\lfloor n/3 \rfloor))$	$U(T_n(\lfloor n/3 \rfloor))$
key recovery ($bikz_1$)	298.87	710.11
randomness recovery ($bikz_2$)	298.87	710.11
plaintext recovery using hints ($bikz_3$)	298.14	712.95
quantum security (bits)	79	188

With the first sample set, we achieve a quantum bit security of 79 and with the second one, we achieve a quantum bit security of 188. We made the SageMath code of our experiments publicly available.⁵

7.4.1 Choice of the Number of Rows

We now discuss the influence of the parameters t , i.e., the number of rows of the Vandermonde matrix chosen to construct $\bar{\mathbf{V}}_\Omega$, on the security of our scheme. This observation also applies to the the original proposal in [HS15].

Increasing t leads to an easier Key Recovery Attack against PASS Encrypt, as the underlying LWE dimension $n-t$ decreases. On the other hand, decreasing t leads to an easier Randomness Recovery Attack against PASS Encrypt, as the underlying LWE dimension t decreases. Hence, choosing $t = \lfloor n/2 \rfloor$ is the optimal choice, as it balances the hardness of both attacks. Our experiments with $t = \lfloor n/3 \rfloor$ (Set A), $t = \lfloor n/2 \rfloor$ (Set B) and $t = \lfloor 2n/3 \rfloor$ (Set C) validate those observations and are summarized in Table 7.2. In both variations (Set A and Set C) the quantum security of PASS Encrypt decreases from 79 to 45.

We emphasize that the observations made above do not apply to the sequence of works on PASS Sign. In the recent publication on the aggregate variant of PASS Sign by Doröz et al. [DHSS20], the parameter t is set to $\lfloor n/3 \rfloor$. Note that there is only the partial Fourier SIS problem with the matrix $\bar{\mathbf{F}}_\Omega \in \mathbb{Z}_q^{t \times n}$ arising in the design of the signature scheme, and not the complement matrix $\bar{\mathbf{F}}_{\Omega^c}$. Hence, decreasing t only makes the corresponding dimension of the LWE instance, that is defined by an instance of partial Fourier SIS, larger and thus the problem harder.

⁵<https://github.com/KatinkaBou/SecurityAnalysisPASSEncrypt>

Table 7.2: Security estimations with different number of rows t for PASS Encrypt. The number of quantum security bits is computed as $\lfloor 0.265 \cdot \min_{j \in [3]}(bikz_j) \rfloor$.

Parameter	Set A	Set B	Set C
ν	2048	2048	2048
n	1024	1024	1024
t	341	512	682
q	12289	12289	12289
key recovery ($bikz_1$)	474.89	298.87	171.82
randomness recovery ($bikz_2$)	171.09	298.87	473.45
plaintext recovery using hints ($bikz_3$)	202.87	298.14	430.49
quantum security (bits)	45	79	45

7.4.2 Comparison

Finally, we provide a comparison between the asymptotic parameters of PASS Encrypt with two other efficient lattice-based PKE schemes. We therefore compute the asymptotic parameters in bits for the secret key sk , the public key pk and the ciphertext \mathbf{c} .

Recall from Section 7.3 that t is the important parameter defining the asymptotic security of PASS Encrypt, motivating our choice to state all parameters with regard to t . In PASS Encrypt the secret key is a ring element sampled from the uniform distribution over $T_n(\lfloor n/3 \rfloor)$. Further, we assume that $n = 2t$, which is the optimal choice as argued in Section 7.4.1. Thus, the bit size of the secret key is $2t \cdot \log_2 3$. The public key is given by $\text{pk} = (\Omega, \overline{\mathbf{V}}_\Omega \text{sk})$ and lies in $\mathbb{Z}_q^t \times \mathbb{Z}_q^t$, requiring $2t \cdot \log_2 q$ bits to transmit it.⁶ Finally, the ciphertext is an element of $\mathbb{Z}_q^t \times \mathbb{Z}_q^{n-t} \times \mathbb{Z}_q^{n-t}$, with $n - t = t$, requiring $3t \cdot \log_2 q$ bits to send it.

We now compare our scheme with two other efficient lattice-based PKE schemes, as illustrated in Table 7.3. The first is the Regev-like PKE scheme based on P-LWE, as presented in [LP11], and the second is the NTRU scheme, as presented in [HPS98]. In [LP11], the secret key and the public key are both ring elements of the ring $R = \mathbb{Z}[x]/\langle x^t + 1 \rangle$, where t is a power of two and the parameter that is determining the asymptotic security. For a better comparison, we assume that the secret is, as in PASS Encrypt, sampled uniformly over $T_t(\lfloor n/3 \rfloor)$. The ciphertext is composed of two ring elements, allowing to encrypt a t -bit message. In [HPS98], the ring $R = \mathbb{Z}[x]/\langle x^t - 1 \rangle$ is used. The secret key is a ring element of small norm. Again, for better comparison, we use the same distribution as in PASS Encrypt. The public key and the ciphertext are elements of R and the schemes allows to encrypt a t -bit message. We note that for simplicity we consider non-optimized versions of the three schemes.

An important characteristic of an PKE scheme is the ratio between the sum of the bit size of its public key and ciphertext and the bit size of the encrypted message. Table 7.3 shows that this ratio is $2.5 \log_2 q$ for PASS Encrypt and thus placing it in between the one of NTRU [HPS98] and the one of the P-LWE-based Regev scheme [LP11]. However, we believe that the algebraic structure of the partial Vandermonde transform $\overline{\mathbf{V}}_\Omega$ used within PASS Encrypt may lead to interesting constructions such as homomorphic commitments or zero-knowledge proofs.

⁶We could further save in storage and bandwidth by only transmitting an index vector in $\{0, 1\}^n$ (instead of the full vector Ω) indicating which row of $\overline{\mathbf{V}}$ is used for the public key.

Table 7.3: Asymptotic parameters in bits for PASS Encrypt, the Regev-like PKE over P-LWE and the NTRU encryption scheme.

Parameter	PASS Encrypt	Regev-like [LP11]	NTRU [HPS98]
$ \mathbf{sk} $	$2t \cdot \log_2 3$	$t \cdot \log_2 3$	$t \cdot \log_2 3$
$ \mathbf{pk} $	$2t \cdot \log_2 q$	$t \cdot \log_2 q$	$t \cdot \log_2 q$
$ \mathbf{c} $	$3t \cdot \log_2 q$	$2t \cdot \log_2 q$	$t \cdot \log_2 q$
$ \mathbf{m} $	$2t$	t	t
$(\mathbf{pk} + \mathbf{c})/ \mathbf{m} $	$2.5 \cdot \log_2 q$	$3 \cdot \log_2 q$	$2 \cdot \log_2 q$

Conclusion and Perspectives

It is now time to conclude the presented contributions of this manuscript and to highlight some short-term and long-term perspectives for future works.

Hardness of Module LWE

In Chapter 2 and Chapter 3, we studied the hardness of Module Learning With Errors (M-LWE). First, we proved in two different ways that its binary secret variant (bin-M-LWE) doesn't become significantly easier as long as we increase the module rank and the noise width accordingly. We were able to generalize this result to any secret whose coefficients lie in a ball of small radius. Those results were then used, among other technical tools, to give a classical worst-case to average-case reduction for M-LWE in Chapter 3. The main drawback of the classical proof is that the rank has to be increased by a factor of at least $\log_2 q$ (inherited from the binary secret reduction before), where the starting modulus q is exponentially large in the ring degree n . This results in a rank that has to be *linear* in n . In practical schemes, however, *small constant* ranks are used. The $\log_2 q$ factor essentially comes from the used Leftover Hash Lemma (LHL), which is a direct generalization of Micciancio's LHL [Mic07] from the ring to the module setting.

One short-term perspective to prove the hardness of bin-M-LWE for smaller ranks (and thus to obtain a classical reduction for ranks below n) is to improve the LHL. For instance, Lin et al. [LWW20] introduced a LHL over rings that doesn't work with a concrete secret distribution, but only depends on its min-entropy. However, it adds a strong condition on the modulus q as it requires it to be inert in the underlying number field. In the popular case of power-of-2 cyclotomics, no such integer q exists. Another LHL over rings is proven by Liu and Wang [LW20], also working with the min-entropy of the secret distribution. They observed that a modulus q which is splitting only in few prime ideals over the number field is more suitable for randomness extraction. As the LHL is an important and widely used tool in cryptography, we think that it is worth investigating all known results to see if we can obtain a general version of it that encompasses all existing algebraic LHL. This would possibly lead to a trade-off between the choice of the modulus q , the secret distribution and the underlying number field, which may finally lead to an improvement of our results as well.

An alternative approach is to avoid the exponentially large modulus q all together, which in turn requires to improve Peikert's classical reduction [Pei09], a seemingly difficult task. Furthermore, the current modulus switching of Section 3.4 heavily depends on the size of the secret. To reduce M-LWE with an exponentially large modulus q to M-LWE with a polynomially large modulus p , we made use of the results of Albrecht and Deo [AD17a]. To be more concrete, we were able to guarantee a polynomial increase in the noise width while restricting the secrets to be binary. In contrast, when using a discrete Gaussian distribution (which depends on q) it is not possible to maintain a polynomial error increase, but only an exponential one. A very interesting question is whether it is possible to improve the results of [AD17a] by making use of the general

algebraic framework and the related tight reductions of Peikert and Pepin [PP19]. If we could get, for instance, interesting results already for Gaussian secret distributions, this would leverage the rank condition for the classical reduction.

From a more general perspective, recent works such as [BD20b, LWW20], including our binary reductions from Chapter 2, have made progress towards understanding the impact that changing the secret distribution has on structured variants of LWE. However, very few research focused on changing the noise distribution of structured variants of LWE. For instance, it is unclear if the reduction from LWE to LWE with binary noise by Micciancio and Peikert [MP13] generalizes to structured variants. Furthermore, it would be interesting to introduce a notion of entropic noise LWE and to study its hardness with respect to the min-entropy provided by the noise distribution, for structured and unstructured flavors of LWE. Such investigation would deepen our understanding of the behavior of structured variants of LWE, such as M-LWE. As those variants are used as hardness assumptions in current lattice-based schemes, we think that this is of utmost importance.

Regarding remaining gaps between properties that have been shown for the unstructured LWE case, but not (yet) for the structured M-LWE case, it may be interesting to look at the sample-preserving search-to-decision reduction by Micciancio and Mol [MM11]. This, however, requires to find a suitable way of using the Fourier transform over number fields, which seems to be a challenging research question.

And finally, as for instance observed in the introduction of Chapter 2, there is a large gap between the parameters resulting from theoretic reductions and the ones obtained by looking at the best known attacks to solve the problem. In practice, one selects the parameters for an (M)LWE-based cryptographic scheme using the latter approach. It is very important to further improve the existing reductions (or the existing attacks) in order to obtain tighter results and hence close (or at least narrow) the gap between theoretical and practical hardness of variants of (M)LWE.

Middle-Product Learning With Rounding

In Chapter 4, we introduced a new hardness assumption, that we named the Middle Product Computational Learning With Rounding (MP-CLWR) problem. It is a combination of two known variants of LWE, the Middle-Product Learning With Errors (MP-LWE) and the Learning With Rounding (LWR) problem, inheriting the security advantage of the first and the simplicity advantage of the latter. We proved that this new assumption is under suitable parameter choices at least as hard as MP-LWE, whose hardness is itself guaranteed by worst-case problems over ideal lattices. We defined the problem in its *computational* variant as for today we don't know how to reduce the hardness of decision MP-LWR from worst-case lattice problems, while maintaining the coefficient-wise rounding and allowing for a polynomially large modulus. As mentioned in the related work section of Chapter 4, Liu and Wang [LW20] addressed this problem for the ring variant of LWR by providing a search-to-decision reduction. However, in order to do so, they defined a new way of rounding. For a short-term perspective, we think that it is important to better understand the relation between their way of rounding and the more standard and widely used coefficient-wise rounding (also used in this thesis). To this end, one has to understand the impact that going from the power basis (defining the coefficient-wise rounding) to some normal integral basis (defining the rounding in [LW20]) has on the reduction parameters. Alternatively, it could be interesting to study if the search-to-decision reduction shown for Ring Learning With Errors (R-LWE) by Roşca et al. [RSW18] can be adapted to the setting of deterministic rounding.

More generally speaking, there are still some interesting open questions with respect to the middle-product variant of LWE. As pointed out by Roşca [Roş20] in the conclusion of her Ph.D

thesis, it would be very intriguing to prove a reduction from MP-LWE back to P-LWE or to M-LWE and hence essentially show their equivalence. As no such equivalence is shown by today, we may assume that MP-LWE is indeed harder than M-LWE.

Furthermore, it would be interesting to see more cryptographic applications for our newly introduced problem. For now, we showed in Chapter 6 how to build a basic encryption scheme whose security is based on MP-CLWR. Intuitively, it seems reasonable to assume that existing cryptographic constructions based on (structured) LWR, such as pseudorandom functions [BPR12, CS19], can be adapted to the middle-product setting. Hence, those schemes would gain with respect to the underlying security assumption (on the expense of a certain loss in efficiency). On the other hand, one could consider existing constructions for MP-LWE such as the identity-based encryption scheme by Lombardi et al. [LVV19] and try to adapt them to the deterministic rounding setting. However, as it is up to today unclear how to construct lattice trapdoors for the LWR setting, this seems to be a difficult task.

Partial Vandermonde Problems and PASS Encrypt

In Chapter 5, we studied several problems related to the partial Vandermonde matrix. More precisely, we looked at the Partial Vandermonde Knapsack problem (PV-Knap) and introduced its dual variant, called Partial Vandermonde Learning With Errors (PV-LWE). We showed that both problems are equivalent to each other. Furthermore, we introduced a leaky variant of PV-Knap that we termed the PASS problem. Its name is motivated by the fact that it serves, together with the decision variant of PV-Knap, as the underlying hardness assumption for an encryption scheme, called PASS Encrypt. This scheme was originally introduced by Hoffstein and Silverman [HS15], but without a security proof. In Chapter 7, we slightly modified PASS Encrypt in order to provide a thorough proof of security assuming the intractability of explicitly stated computational problems. Moreover, we proposed a refined analysis of its practical security and concrete sample parameters. By providing additional information on the underlying secret of a PV-Knap instance, the PASS problem seems to be easier than the original PV-Knap. As a short-term perspective, it would be interesting to see if we can build an encryption scheme whose security is based only on PV-Knap (or equivalently PV-LWE). In this direction, in an ongoing work with Amin Sakzad and Ron Steinfeld, we define an encryption scheme using PV-LWE in the spirit of Regev's encryption scheme for LWE [Reg05]. The homomorphic structure of PV-Knap, that comes naturally with the partial Vandermonde matrix, may allow for more efficient and/or more functional constructions. Furthermore, it may be useful to design more advanced schemes, like zero-knowledge proofs for valid commitment openings or for group signatures.

More generally, it is important to continue investigating the hardness of partial Vandermonde problems. As it is a very recent source of hardness assumptions, it is crucial to deepen our understanding how difficult it is to solve those problems. The amount of resources spent to attack a cryptographic system and to solve the underlying hardness assumption is an important characteristic to gain confidence in the proposed scheme.

One possibility is to connect the problems PV-Knap or PV-LWE with worst-case lattice problems, as this is the case for standard Knapsack and LWE [Ajt96, Reg05]. However, the fact that decision PV-LWE becomes easy to solve if the underlying secret is small, could be seen as a warning.

Alternatively, one could try to connect the NTRU problem with the partial Vandermonde problems. Even though, the standard search variant of NTRU does up to today not possess a connection to worst-case (structured) lattice problems, it has been studied for more than 20 years now and seems to be a reasonable hardness assumption in practice.

To tackle the question from another starting point, one could try to design attacks that

take into account the algebraic structure of the Vandermonde matrix to show that the partial Vandermonde problems are possibly easier than standard LWE or SIS problems. A hint towards this reasoning is that the homogeneous variant of PV-Knap, that we called PV-SIS, essentially defines average-case instances of the Ideal Shortest Vector Problem (Id-SVP), which has shown to be an easier problem than the more general Module Shortest Vector Problem (Mod-SVP) for modules of rank at least 2 [CDW17, PHS19, BR20]. Additionally, the attacks by Pan et al. [PXWC21] apply to the specific setting of PV-SIS and thus add restrictions on the modulus q , as it has to be fully splitting in the underlying number field.

General Perspectives for Structured Variants of LWE

In the last 10 years, plenty of different structured variants of LWE have been introduced. The overall goal of this Ph.D thesis was to investigate different choices of structured variants of LWE in order to obtain scientific criteria which concrete flavor to choose for cryptographic schemes. In this work we covered the variants over rings of polynomials (P-LWE), over rings of integers (R-LWE), over modules (M-LWE), using the middle-product (MP-LWE) or deterministic rounding (R-LWR, MP-LWR). Every problem possesses a formulation as a search or a decision problem, but not in every case they are connected via a search-to-decision reduction. Furthermore, we can vary the space of the secret, leading to the problem in the dual or the primal version, or modify the secret and noise distribution, independently. All this has an impact on the hardness of the problem. So, which problem to choose for designing a lattice-based cryptographic scheme? The honest but maybe frustrating answer is that it depends on the concrete scheme, the required security guarantees and the available resources. Regarding the different variants, Peikert and Pepin [PP19] were able to generalize all existing structured variants and show tight reductions starting from R-LWE. In particular, they tightly reduce the problem R-LWE over a number field of degree nd to the problem M-LWE over a number field of degree n and rank d . The authors commented their results as follows [PP19, Abstract]:

A main message of our work is that it is straightforward to use the hardness of the original R-LWE problem as a foundation for the hardness of all other algebraic LWE problems defined over number fields, via simple and rather tight reductions.

Another way to interpret their work is to see M-LWE as the more reliable hardness assumption as it generalizes R-LWE but is at the same time tightly connected to it. In the status report of the NIST's standardization process, they phrased it in the following way [AASA⁺20, Sec. 3.2]:

Recent theoretical work has placed M-LWE on stronger footing by providing a very tight reduction from R-LWE to M-LWE [PP19].

As there are (non-tight) reductions in the opposite direction, showing that M-LWE reduces to R-LWE [AD17a], both problems are essentially equivalent. Using the (non-tight) equivalence between M-LWE and SIVP over module lattices [LS15], we can deduce that all three problems are more or less equally hard. For R-LWE, we only know a worst-case to average-case reduction from SIVP over ideal lattices [LPR10], but not in the other direction. At the same time, there are algorithms that can solve SVP over ideal lattices in sub-exponential time [CDW17, PHS19, BR20], which is up to date not known for SVP over modules (of rank larger than 1). So it seems that there is a difficulty gap between the set of problems {Id-SVP, Id-SIVP} and the set of problems {Mod-SIVP, M-LWE, R-LWE}. Better understanding this potential difficulty gap is one of the major open tasks for the theoretical foundations of lattice-based cryptography for the near future.

More generally, it is possible that Conjecture 2 from the introduction is much stronger than Conjecture 1 in the sense that lattice problems over module lattices may be easier than over general Euclidean lattices. Here again, we still need to invest more research resources to better understand their relation. One very interesting open problem is to tightly connect both sets of problems, which may be even impossible.

From a global perspective, we have no doubt that structured lattice problems will play an important role in public key cryptography after the advent of quantum computers and agree with NIST's statement [AASA⁺20, Sec. 2.3]:

In NIST's current view, these structured lattice schemes appear to be the most promising general-purpose algorithms for public-key encryption/KEM and digital signature schemes.

Nonetheless, it will be interesting to see which structured variants will make the race in the future. Probably, it will mostly depend on their suitability for building concrete and advanced cryptographic schemes.

Bibliography

- [AA16] J. Alperin-Sheriff and D. Apon. Dimension-preserving reductions from LWE to LWR. *IACR Cryptology ePrint Archive*, 2016:589, 2016.
- [AASA⁺20] G. Alagic, J. Alperin-Sheriff, D. Apon, D. Cooper, Q. Dang, J. Kelsey, Y.-K. Liu, C. Miller, D. Moody, R. Peralta, R. Perlner, A. Robinson, and D. Smith-Tone. Status report on the second round of the nist post-quantum cryptography standardization process. *U.S. Department of Commerce, NIST*, 2020. <https://doi.org/10.6028/NIST.IR.8309>, last accessed on 29.06.2021.
- [ACD⁺18] M. R. Albrecht, B. R. Curtis, A. Deo, A. Davidson, R. Player, E. W. Postlethwaite, F. Virdia, and T. Wunderer. Estimate all the {LWE, NTRU} schemes! In *Security and Cryptography for Networks - 11th International Conference, SCN 2018, Amalfi, Italy, September 5-7, 2018, Proceedings*, pages 351–367, 2018.
- [ACF⁺15] M. R. Albrecht, C. Cid, J.-C. Faugère, R. Fitzpatrick, and L. Perret. Algebraic algorithms for LWE problems. *ACM Commun. Comput. Algebra*, 49(2):62, 2015.
- [ACPS09] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618. Springer, 2009.
- [AD97] M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 284–293. ACM, 1997.
- [AD17a] M. R. Albrecht and A. Deo. Large modulus ring-lwe \geq module-lwe. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I*, pages 267–296, 2017.
- [AD17b] M. R. Albrecht and A. Deo. Large modulus ring-lwe \geq module-lwe. *IACR Cryptol. ePrint Arch.*, 2017:612, 2017.
- [ADPS16] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. Post-quantum key exchange - A new hope. In *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016.*, pages 327–343, 2016.

- [AG11] S. Arora and R. Ge. New algorithms for learning in presence of errors. In *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part I*, pages 403–415, 2011.
- [Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 99–108. ACM, 1996.
- [AKPW13] J. Alwen, S. Krenn, K. Pietrzak, and D. Wichs. Learning with rounding, revisited - new reduction, properties and applications. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, pages 57–74, 2013.
- [APS15] M. R. Albrecht, R. Player, and S. Scott. On the concrete hardness of learning with errors. *J. Mathematical Cryptology*, 9(3):169–203, 2015.
- [Bab85] L. Babai. On Lovász’ lattice reduction and the nearest lattice point problem. In *STACS 85 (Saarbrücken, 1985)*, volume 182 of *Lecture Notes in Comput. Sci.*, pages 13–20. Springer, Berlin, 1985.
- [Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Math. Ann.*, 296(4):625–635, 1993.
- [BBD⁺19] S. Bai, K. Boudgoust, D. Das, A. Roux-Langlois, W. Wen, and Z. Zhang. Middle-product learning with rounding problem and its applications. In *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part I*, pages 55–81, 2019.
- [BBF⁺19] H. Baan, S. Bhattacharya, S. R. Fluhrer, Ó. García-Morchón, T. Laarhoven, R. Ritman, M.-J. O. Saarinen, L. Tolhuizen, and Z. Zhang. Round5: Compact and fast post-quantum public-key encryption. In *Post-Quantum Cryptography - 10th International Conference, PQCrypto 2019, Chongqing, China, May 8-10, 2019 Revised Selected Papers*, pages 83–102, 2019.
- [BD20a] Z. Brakerski and N. Döttling. Hardness of LWE on general entropic distributions. In *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part II*, volume 12106 of *Lecture Notes in Computer Science*, pages 551–575. Springer, 2020.
- [BD20b] Z. Brakerski and N. Döttling. Lossiness and entropic hardness for ring-lwe. In *Theory of Cryptography - 18th International Conference, TCC 2020, Durham, NC, USA, November 16-19, 2020, Proceedings, Part I*, volume 12550 of *Lecture Notes in Computer Science*, pages 1–27. Springer, 2020.
- [BDH⁺20] S. Bai, D. Das, R. Hiromasa, M. Rosca, A. Sakzad, D. Stehlé, R. Steinfeld, and Z. Zhang. Mpsign: A signature from small-secret middle-product learning with errors. In *Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part II*, volume 12111 of *Lecture Notes in Computer Science*, pages 66–93. Springer, 2020.

- [BDJ06] W. Bryc, A. Dembo, and T. Jiang. Spectral measure of large random Hankel, Markov and Toeplitz matrices. *Ann. Probab.*, 34(1):1–38, 2006.
- [BDK⁺18] J. W. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS - kyber: A cca-secure module-lattice-based KEM. In *2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018, London, United Kingdom, April 24-26, 2018*, pages 353–367, 2018.
- [BGM⁺16] A. Bogdanov, S. Guo, D. Masny, S. Richelson, and A. Rosen. On the hardness of learning with rounding over small modulus. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, pages 209–224, 2016.
- [BGPW16] J. Buchmann, F. Göpfert, R. Player, and T. Wunderer. On the hardness of LWE with binary error: Revisiting the hybrid lattice-reduction and meet-in-the-middle attack. In *Progress in Cryptology - AFRICACRYPT 2016 - 8th International Conference on Cryptology in Africa, Fes, Morocco, April 13-15, 2016, Proceedings*, volume 9646 of *Lecture Notes in Computer Science*, pages 24–43. Springer, 2016.
- [BGV12] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 309–325, 2012.
- [BHLY16] L. Groot Bruinderink, A. Hülsing, T. Lange, and Y. Yarom. Flush, gauss, and reload - A cache attack on the BLISS lattice-based signature scheme. In *Cryptographic Hardware and Embedded Systems - CHES 2016 - 18th International Conference, Santa Barbara, CA, USA, August 17-19, 2016, Proceedings*, pages 323–345, 2016.
- [BJRW20] K. Boudgoust, C. Jeudy, A. Roux-Langlois, and W. Wen. Towards classical hardness of module-lwe: The linear rank case. In *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part II*, volume 12492 of *Lecture Notes in Computer Science*, pages 289–317. Springer, 2020.
- [BJRW21] K. Boudgoust, C. Jeudy, A. Roux-Langlois, and W. Wen. On the hardness of module-lwe with binary secret. In *Topics in Cryptology - CT-RSA 2021 - Cryptographers' Track at the RSA Conference 2021, Virtual Event, May 17-20, 2021, Proceedings*, volume 12704 of *Lecture Notes in Computer Science*, pages 503–526. Springer, 2021.
- [BLL⁺15] S. Bai, A. Langlois, T. Lepoint, D. Stehlé, and R. Steinfeld. Improved security proofs in lattice-based cryptography: Using the rényi divergence rather than the statistical distance. In *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part I*, pages 3–24, 2015.
- [BLP⁺13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 575–584, 2013.
- [BLR⁺18] S. Bai, T. Lepoint, A. Roux-Langlois, A. Sakzad, D. Stehlé, and R. Steinfeld. Improved security proofs in lattice-based cryptography: Using the rényi divergence rather than the statistical distance. *J. Cryptol.*, 31(2):610–640, 2018.

- [BPR11] A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. *IACR Cryptology ePrint Archive*, 2011:401, 2011.
- [BPR12] A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 719–737, 2012.
- [BR20] O. Bernard and A. Roux-Langlois. Twisted-phs: Using the product formula to solve approx-svp in ideal lattices. In *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part II*, volume 12492 of *Lecture Notes in Computer Science*, pages 349–380. Springer, 2020.
- [BV14] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM J. Comput.*, 43(2):831–871, 2014.
- [CCLS20] H. Chen, L. Chua, K. E. Lauter, and Y. Song. On the concrete security of LWE with small secret. *IACR Cryptol. ePrint Arch.*, 2020:539, 2020.
- [CDH⁺] C. Chen, O. Danba, J. Hoffstein, A. Hülsing, J. Rijneveld, T. Saito, J. M. Schanck, P. Schwabe, W. Whyte, K. Xagawa, T. Yamakawa, and Z. Zhang. Ntru - a submission to the nist post-quantum standardization effort. <https://ntru.org>, last accessed on 08.07.2021.
- [CDPR16] R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, pages 559–585, 2016.
- [CDW17] R. Cramer, L. Ducas, and B. Wesolowski. Short stickelberger class relations and application to ideal-svp. In *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part I*, pages 324–348, 2017.
- [CGGI16] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène. Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In *Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I*, pages 3–33, 2016.
- [CIV16] W. Castryck, I. Iliashenko, and F. Vercauteren. Provably weak instances of ring-lwe revisited. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part I*, pages 147–167, 2016.
- [Con] K. Conrad. The different ideal. <https://kconrad.math.uconn.edu/blurbs/gradnumthy/different.pdf>, last accessed on 08.07.2021.

- [CS19] C. Chuengsatiansup and D. Stehlé. Towards practical ggm-based PRF from (module-)learning-with-rounding. In *Selected Areas in Cryptography - SAC 2019 - 26th International Conference, Waterloo, ON, Canada, August 12-16, 2019, Revised Selected Papers*, volume 11959 of *Lecture Notes in Computer Science*, pages 693–713. Springer, 2019.
- [CZZ18] L. Chen, Z. Zhang, and Z. Zhang. On the hardness of the computational ring-lwr problem and its applications. In *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part I*, pages 435–464, 2018.
- [DAZ19] D. Das, M. Ho Au, and Z. Zhang. Ring signatures based on middle-product learning with errors problems. In *Progress in Cryptology - AFRICACRYPT 2019 - 11th International Conference on Cryptology in Africa, Rabat, Morocco, July 9-11, 2019, Proceedings*, volume 11627 of *Lecture Notes in Computer Science*, pages 139–156. Springer, 2019.
- [DB15] C. Du and G. Bai. Towards efficient discrete gaussian sampling for lattice-based cryptography. In *25th International Conference on Field Programmable Logic and Applications, FPL 2015, London, United Kingdom, September 2-4, 2015*, pages 1–6, 2015.
- [DDGR20] D. Dachman-Soled, L. Ducas, H. Gong, and M. Rossi. LWE with side information: Attacks and concrete security estimation. In *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 329–358. Springer, 2020.
- [DH76] W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Trans. Inf. Theory*, 22(6):644–654, 1976.
- [DHSS20] Y. Doröz, J. Hoffstein, J. H. Silverman, and B. Sunar. MMSAT: A scheme for multimessage multiuser signature aggregation. *IACR Cryptol. ePrint Arch.*, 2020:520, 2020.
- [DKL⁺18] L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.
- [DKRV18] J.-P. D’Anvers, A. Karmakar, S. Sinha Roy, and F. Vercauteren. Saber: Module-lwr based key exchange, cpa-secure encryption and cca-secure KEM. In *Progress in Cryptology - AFRICACRYPT 2018 - 10th International Conference on Cryptology in Africa, Marrakesh, Morocco, May 7-9, 2018, Proceedings*, pages 282–305, 2018.
- [DM15] L. Ducas and D. Micciancio. FHEW: bootstrapping homomorphic encryption in less than a second. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, pages 617–640, 2015.
- [DMR18] L. Devroye, A. Mehrabian, and T. Reddad. The total variation distance between high-dimensional gaussians, 2018.

- [DR00] J. Daemen and V. Rijmen. Rijndael for AES. In *The Third Advanced Encryption Standard Candidate Conference, April 13-14, 2000, New York, New York, USA*, pages 343–348. National Institute of Standards and Technology, 2000.
- [DXL12] J. Ding, X. Xie, and X. Lin. A simple provably secure key exchange scheme based on the learning with errors problem. *IACR Cryptology ePrint Archive*, 2012:688, 2012.
- [Gam] Jay Gambetta. Ibm’s roadmap for scaling quantum technology. <https://www.ibm.com/blogs/research/2020/09/ibm-quantum-roadmap/>, last accessed on 08.07.2021.
- [Gam84] T. El Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *Advances in Cryptology, Proceedings of CRYPTO '84, Santa Barbara, California, USA, August 19-22, 1984, Proceedings*, volume 196 of *Lecture Notes in Computer Science*, pages 10–18. Springer, 1984.
- [GE21] C. Gidney and M. Ekerå. How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits. *Quantum*, 5:433, 2021.
- [GHPS12] C. Gentry, S. Halevi, C. Peikert, and N. P. Smart. Ring switching in bgv-style homomorphic encryption. In *Security and Cryptography for Networks - 8th International Conference, SCN 2012, Amalfi, Italy, September 5-7, 2012. Proceedings*, volume 7485 of *Lecture Notes in Computer Science*, pages 19–37. Springer, 2012.
- [GKPV10] S. Goldwasser, Y. T. Kalai, C. Peikert, and V. Vaikuntanathan. Robustness of the learning with errors assumption. In *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings*, pages 230–240. Tsinghua University Press, 2010.
- [GOPS13] T. Güneysu, T. Oder, T. Pöppelmann, and P. Schwabe. Software speed records for lattice-based signatures. In *Post-Quantum Cryptography - 5th International Workshop, PQCrypto 2013, Limoges, France, June 4-7, 2013. Proceedings*, volume 7932 of *Lecture Notes in Computer Science*, pages 67–82. Springer, 2013.
- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 197–206. ACM, 2008.
- [Gro97] L. K. Grover. Quantum computers can search arbitrarily large databases by a single query. *Physical review letters*, 79(23):4709, 1997.
- [GS21] É. Gouzien and N. Sangouard. Factoring 2048 rsa integers in 177 days with 13436 qubits and a multimode memory, 2021.
- [HH18] T. Hurley and D. Hurley. Coding theory: the unit-derived methodology. *IJICoT*, 5(1):55–80, 2018.
- [HILL99] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [HPS98] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In *Algorithmic Number Theory, Third International Symposium, ANTS-III, Portland, Oregon, USA, June 21-25, 1998, Proceedings*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998.

- [HPS⁺14] J. Hoffstein, J. Pipher, J. M. Schanck, J. H. Silverman, and W. Whyte. Practical signatures from the partial fourier recovery problem. In *Applied Cryptography and Network Security - 12th International Conference, ACNS 2014, Lausanne, Switzerland, June 10-13, 2014. Proceedings*, pages 476–493, 2014.
- [HS15] J. Hoffstein and J. H. Silverman. Pass-encrypt: a public key cryptosystem based on partial evaluation of polynomials. *Des. Codes Cryptogr.*, 77(2-3):541–552, 2015.
- [KF15] P. Kirchner and P.-A. Fouque. An improved BKW algorithm for LWE with applications to cryptography and lattices. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, volume 9215 of *Lecture Notes in Computer Science*, pages 43–62. Springer, 2015.
- [KL96] E. Kaltofen and A. Lobo. On rank properties of toeplitz matrices over finite fields. In *ISSAC*, volume 96, pages 241–249, 1996.
- [KL14] J. Katz and Y. Lindell. *Introduction to Modern Cryptography, Second Edition*. CRC Press, 2014.
- [Laa15] T. Laarhoven. Search problems in cryptography, 2015. <http://www.thijs.com/docs/phd-final.pdf>, last accessed on 08.07.2021.
- [LDSP20] H. Q. Le, D. H. Duong, W. Susilo, and J. Pieprzyk. Trapdoor delegation and HIBE from middle-product LWE in standard model. In *Applied Cryptography and Network Security - 18th International Conference, ACNS 2020, Rome, Italy, October 19-22, 2020, Proceedings, Part I*, volume 12146 of *Lecture Notes in Computer Science*, pages 130–149. Springer, 2020.
- [LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261(4):515–534, 1982.
- [LM00] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Ann. Statist.*, 28(5):1302–1338, 2000.
- [LNS21] V. Lyubashevsky, N. K. Nguyen, and G. Seiler. Shorter lattice-based zero-knowledge proofs via one-time commitments. In Juan A. Garay, editor, *Public-Key Cryptography - PKC 2021 - 24th IACR International Conference on Practice and Theory of Public Key Cryptography, Virtual Event, May 10-13, 2021, Proceedings, Part I*, volume 12710 of *Lecture Notes in Computer Science*, pages 215–241. Springer, 2021.
- [LP11] R. Lindner and C. Peikert. Better key sizes (and attacks) for lwe-based encryption. In *Topics in Cryptology - CT-RSA 2011 - The Cryptographers’ Track at the RSA Conference 2011, San Francisco, CA, USA, February 14-18, 2011. Proceedings*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339. Springer, 2011.
- [LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings*, pages 1–23, 2010.
- [LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43:1–43:35, 2013.

- [LS15] A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.*, 75(3):565–599, 2015.
- [LS18] V. Lyubashevsky and G. Seiler. Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 204–224. Springer, 2018.
- [LSS14] A. Langlois, D. Stehlé, and R. Steinfeld. Gghlite: More efficient multilinear maps from ideal lattices. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 239–256. Springer, 2014.
- [LVV19] A. Lombardi, V. Vaikuntanathan, and T. Duong Vuong. Lattice trapdoors and IBE from middle-product LWE. In *Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part I*, volume 11891 of *Lecture Notes in Computer Science*, pages 24–54. Springer, 2019.
- [LW20] F.-H. Liu and Z. Wang. Rounding in the rings. In *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 296–326. Springer, 2020.
- [LWW20] H. Lin, Y. Wang, and M. Wang. Hardness of module-lwe and ring-lwe on general entropic distributions. *IACR Cryptol. ePrint Arch.*, 2020:1238, 2020.
- [Lyu09] V. Lyubashevsky. Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *Lecture Notes in Computer Science*, pages 598–616. Springer, 2009.
- [Lyu12] V. Lyubashevsky. Lattice signatures without trapdoors. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755. Springer, 2012.
- [LZA18] X. Lu, Z. Zhang, and M. H. Au. Practical signatures from the partial fourier recovery problem revisited: A provably-secure and gaussian-distributed construction. In *Information Security and Privacy - 23rd Australasian Conference, ACISP 2018, Wollongong, NSW, Australia, July 11-13, 2018, Proceedings*, pages 813–820, 2018.
- [Mer78] R. C. Merkle. Secure communications over insecure channels. *Commun. ACM*, 21(4):294–299, 1978.
- [Mic07] D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complex.*, 16(4):365–411, 2007.

- [Mic18] D. Micciancio. On the hardness of learning with errors with binary secrets. *Theory of Computing*, 14(1):1–17, 2018.
- [MM11] D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 465–484. Springer, 2011.
- [MP12] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 700–718, 2012.
- [MP13] D. Micciancio and C. Peikert. Hardness of SIS and LWE with small parameters. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 2013.
- [MR07] D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.
- [MR09] D. Micciancio, , and O. Regev. Lattice-based cryptography. In *Post-Quantum Cryptography*, pages 147–191. Springer, 2009.
- [NIS] NIST. Post-quantum cryptography standardization. <https://csrc.nist.gov/Projects/Post-Quantum-Cryptography>, last accessed on 08.07.2021.
- [Pei08] C. Peikert. Limits on the hardness of lattice problems in l_p norms. *Comput. Complex.*, 17(2):300–351, 2008.
- [Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 333–342, 2009.
- [Pei10] C. Peikert. An efficient and parallel gaussian sampler for lattices. In *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, pages 80–97, 2010.
- [Pei14] C. Peikert. Lattice cryptography for the internet. In *Post-Quantum Cryptography - 6th International Workshop, PQCrypto 2014, Waterloo, ON, Canada, October 1-3, 2014. Proceedings*, pages 197–219, 2014.
- [Pei16a] C. Peikert. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.*, 10(4):283–424, 2016.
- [Pei16b] C. Peikert. How (not) to instantiate ring-lwe. In *Security and Cryptography for Networks - 10th International Conference, SCN 2016, Amalfi, Italy, August 31 - September 2, 2016, Proceedings*, pages 411–430, 2016.
- [Pes16] P. Pessl. Analyzing the shuffling side-channel countermeasure for lattice-based signatures. In *Progress in Cryptology - INDOCRYPT 2016 - 17th International Conference on Cryptology in India, Kolkata, India, December 11-14, 2016, Proceedings*, pages 153–170, 2016.

- [PHS19] A. Pellet-Mary, G. Hanrot, and D. Stehlé. Approx-svp in ideal lattices with pre-processing. In *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part II*, pages 685–716, 2019.
- [PP19] Chris Peikert and Zachary Pepin. Algebraically structured lwe, revisited. In *TCC (1)*, volume 11891 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2019.
- [PR07] C. Peikert and A. Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 478–487, 2007.
- [PRS17a] C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of ring-lwe for any ring and modulus. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 461–473, 2017.
- [PRS17b] C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of ring-lwe for any ring and modulus. *IACR Cryptol. ePrint Arch.*, 2017:258, 2017.
- [PS19] C. Peikert and S. Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 89–114. Springer, 2019.
- [PXWC21] Y. Pan, J. Xu, N. Wadleigh, and Q. Cheng. On the ideal shortest vector problem over random rational primes. In *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part I*, volume 12696 of *Lecture Notes in Computer Science*, pages 559–583. Springer, 2021.
- [R61] A. Rényi. On measures of entropy and information. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. I*, pages 547–561. Univ. California Press, Berkeley, Calif., 1961.
- [Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93, 2005.
- [Reg09] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, 2009.
- [Roş20] M. Roşca. *On algebraic variants of Learning With Errors. (Sur des variantes algébriques du problème Learning With Errors)*. PhD thesis, University of Lyon, France, 2020.
- [Rot06] R. M. Roth. *Introduction to coding theory*. Cambridge University Press, 2006.
- [RSA78] R. L. Rivest, A. Shamir, and L. M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.

- [RSSS17] M. Roşca, A. Sakzad, D. Stehlé, and R. Steinfeld. Middle-product learning with errors. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part III*, pages 283–297, 2017.
- [RSW18] M. Roşca, D. Stehlé, and A. Wallet. On the ring-lwe and polynomial-lwe problems. In *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, pages 146–173, 2018.
- [Saa18] M.-J. O. Saarinen. Arithmetic coding and blinding countermeasures for lattice signatures - engineering a side-channel resistant post-quantum signature scheme with compact signatures. *J. Cryptographic Engineering*, 8(1):71–84, 2018.
- [Sch87] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.
- [SD16] N. Stephens-Davidowitz. Dimension-preserving reductions between lattice problems, 2016. <http://noahsd.com/latticeproblems.pdf>, last accessed on 08.07.2021.
- [SE94] C.-P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Math. Program.*, 66:181–199, 1994.
- [Sho97] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.
- [SS11] D. Stehlé and R. Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, volume 6632 of *Lecture Notes in Computer Science*, pages 27–47. Springer, 2011.
- [SS13] D. Stehlé and R. Steinfeld. Making ntruencrypt and ntrusign as secure as standard worst-case problems over ideal lattices. *IACR Cryptol. ePrint Arch.*, 2013:4, 2013.
- [SSTX09] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, pages 617–635, 2009.
- [SSZ17] R. Steinfeld, A. Sakzad, and R. K. Zhao. Titanium: Proposal for a nist post-quantum public-key encryption and kem standard, 2017. <http://users.monash.edu.au/~rste/Titanium.html>, last accessed on 08.07.2021.
- [SSZ19] R. Steinfeld, A. Sakzad, and R. K. Zhao. Practical mp-lwe-based encryption balancing security-risk versus efficiency. *Des. Codes Cryptogr.*, 87(12):2847–2884, 2019.
- [STA20] C. Sun, M. Tibouchi, and M. Abe. Revisiting the hardness of binary error LWE. In *Information Security and Privacy - 25th Australasian Conference, ACISP 2020, Perth, WA, Australia, November 30 - December 2, 2020, Proceedings*, volume 12248 of *Lecture Notes in Computer Science*, pages 425–444. Springer, 2020.
- [vEH14] T. van Erven and P. Harremoës. Rényi divergence and kullback-leibler divergence. *IEEE Trans. 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.
- [YXW17] Y. Yu, G. Xu, and X. Wang. Provably secure ntruencrypt over more general cyclotomic rings. *IACR Cryptology ePrint Archive*, 2017:304, 2017.

Titre : Difficulté théorique des variantes algébriques du problème Learning With Errors

Mots-clés : cryptographie à base de réseaux, apprentissage avec erreurs, variantes structurées, secret binaire, difficulté classique, produit intermédiaire, matrice de Vandermonde

Résumé : Cette thèse de doctorat porte principalement sur le problème d'apprentissage avec erreurs, appelé Learning With Errors (LWE). Il s'agit d'une composante essentielle de la cryptographie à base de réseaux, qui fait partie des candidats les plus prometteurs pour remplacer les protocoles cryptographiques actuels lorsque des ordinateurs quantiques à grande échelle seront disponibles.

Dans cette thèse, nous étudions la difficulté théoriques des variantes algébriquement structurées de LWE qui sont utilisées dans des protocoles efficaces. D'abord, nous prouvons que le problème Module Learning With Errors (M-LWE) ne devient pas significativement plus facile à résoudre, même si le secret sous-jacent est remplacé par un vecteur binaire. Ensuite, nous présentons une réduction classique au problème M-LWE, ce qui renforce notre confiance dans sa valeur pour la cryptographie. De plus, nous définissons une nouvelle hypo-

thèse de difficulté, le problème MP-CLWR (Middle-Product Computational Learning With Rounding), qui hérite des avantages de deux variantes existantes de LWE. Enfin, nous étudions des problèmes liés à la matrice partielle de Vandermonde. Il s'agit d'une source récente d'hypothèses de difficulté pour la cryptographie à base de réseaux et son étude rigoureuse est primordiale pour gagner en fiabilité. Finalement, nous montrons que les nouvelles hypothèses de difficulté introduites auparavant servent à la construction de schémas de chiffrement à clé publique efficaces. D'une part, nous concevons un nouveau schéma de chiffrement, dont la sécurité est assurée par la difficulté du problème MP-CLWR. D'autre part, nous modifions un schéma de chiffrement existant pour lui fournir une preuve de sécurité basée sur deux problèmes de Vandermonde partiels explicitement énoncés.

Title: Theoretical Hardness of Algebraically Structured Learning With Errors

Keywords: lattice-based cryptography, learning with errors, structured variants, binary secrets, classical hardness, middle-product, Vandermonde matrix

Abstract: The main focus of this Ph.D thesis lies on the computational problem Learning With Errors (LWE). It is a core building block of lattice-based cryptography, which itself is among the most promising candidates to replace current cryptographic protocols once large-scale quantum computers may be available.

The contributions of the present work are separated into two different parts. First, we study the hardness of structured variants of LWE. To this end, we show that under suitable parameter choices the Module Learning With Errors (M-LWE) problem doesn't become significantly easier to solve even if the underlying secret is replaced by a binary vector. Furthermore, we provide a classical hardness reduction for M-LWE, which further strengthens our confidence in its suitability for cryptography. Ad-

ditionally, we define a new hardness assumption, the Middle-Product Computational Learning With Rounding (MP-CLWR) problem, which inherits the advantages of two existing LWE variants. Finally, we study problems related to the partial Vandermonde matrix. This is a recent source of hardness assumptions for lattice-based cryptography and its rigorous study is important to gain trust in it. In the second part of this manuscript, we show that the new hardness assumptions we introduced before serve for the construction of efficient public-key encryption. On the one hand, we design a new encryption scheme, whose security is provably based on the MP-CLWR problem. On the other hand, we modify an existing encryption scheme, called PASS Encrypt, to provide it with a security proof based on two explicitly stated partial Vandermonde problems.