

On algebraic variants of Learning With Errors

Georgiana-Miruna Rosca

▶ To cite this version:

Georgiana-Miruna Rosca. On algebraic variants of Learning With Errors. Cryptography and Security [cs.CR]. Université de Lyon, 2020. English. NNT: 2020LYSEN063. tel-03085029

HAL Id: tel-03085029 https://theses.hal.science/tel-03085029

Submitted on 21 Dec 2020 $\,$

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.



Numéro National de Thèse : 2020LYSEN063

THÈSE DE DOCTORAT DE L'UNIVERSITÉ DE LYON opérée au sein de l'École Normale Supérieure de Lyon

École Doctorale N°512 École Doctorale en Informatique et Mathématiques de Lyon

Spécialité de doctorat : Informatique

Soutenue publiquement le 17/11/2020, par : Georgiana-Miruna Roșca

On algebraic variants of Learning With Errors

Sur des variantes algébriques du problème Learning With Errors

Devant le jury composé de :

LYUBASHEVSKY Vadim, Chercheur, IBM Zürich (Suisse)	Rapporteur
VERCAUTEREN Frederik, Professeur Associé, KU Leuven (Belgique)	Rapporteur
GABORIT Philippe, Professeur des Universités, Université de Limoges	Examinateur
ROUX-LANGLOIS Adeline, Chargée de recherche, Université Rennes 1, CNRS et IRISA	Examinatrice
VALLÉE Brigitte, Directrice de Recherche émérite, CNRS et Université de Caen	Examinatrice

STEHLÉ Damien, Professeur des Universités, ENS de Lyon

Directeur de thèse

Résumé

La cryptographie à base de réseaux euclidiens repose en grande partie sur l'utilisation du problème Learning With Errors (LWE) comme fondation de sécurité. Ce problème est au moins aussi difficile que les problèmes standards portant sur les réseaux, mais les primitives cryptographiques qui l'utilisent sont inefficaces en termes de consommation en temps et en espace. Les problèmes Polynomial Learning With Errors (PLWE), dual Ring Learning With Errors (dual-RLWE) et primal Ring Learning With Errors (primal-RLWE) sont trois variantes de LWE qui utilisent des stuctures algébriques supplémentaires afin de pallier les inconvénients ci-dessus. Le problème PLWE est paramétré par un polynôme f, alors que dual-RLWE et primal-RLWE sont définis à l'aide de l'anneau d'entiers d'un corps de nombres. Ces problèmes, dits algébriques, sont eux-mêmes au moins aussi difficiles que des problèmes standards portant sur les réseaux, mais, dans leur cas, les réseaux impliqués appartiennent à des classes restreintes. Dans cette thèse, nous nous intéressons aux liens entre les variantes algébriques de LWE.

Tout d'abord, nous montrons que pour une vaste classe de polynômes de définition, il existe des réductions (non-uniformes) entre dual-RLWE, primal-RLWE et PLWE pour lesquelles l'amplification des paramètres peut être contrôlée. Ces résultats peuvent être interprétés comme une indication forte de l'équivalence calculatoire de ces problèmes.

Ensuite, nous introduisons une nouvelle variante algébrique de LWE, *Middle-Product Learning With Errors* (MP-LWE). On montre que ce problème est au moins aussi difficile que PLWE pour beaucoup de polynômes de définition f. Par conséquent, un système cryptographique reposant sur MP-LWE reste sûr aussi longtemps qu'une de ces instances de PLWE reste difficile à résoudre.

Enfin, nous montrons la pertinence cryptographique de MP-LWE en proposant un protocole de chiffrement asymétrique et une signature digitale dont la sécurité repose sur la difficulté présumée de MP-LWE.

Abstract

Lattice-based cryptography relies in great parts on the use of the Learning With Errors (LWE) problem as hardness foundation. This problem is at least as hard as standard worst-case lattice problems, but the primitives based on it usually have big key sizes and slow algorithms. Polynomial Learning With Errors (PLWE), dual Ring Learning With Errors (dual-RLWE) and primal Ring Learning With Errors (primal-RLWE) are variants of LWE which make use of extra algebraic structures in order to fix the above drawbacks. The PLWE problem is parameterized by a polynomial f, while dual-RLWE and primal-RLWE are defined using the ring of integers of a number field. These problems, which we call *algebraic*, also enjoy reductions from worst-case lattice problems, but in their case, the lattices involved belong to diverse restricted classes.

In this thesis, we study relationships between algebraic variants of LWE.

We first show that for many defining polynomials, there exist (non-uniform) reductions between dual-RLWE, primal-RLWE and PLWE that incur limited parameter losses. These results could be interpreted as a strong evidence that these problems are qualitatively equivalent.

Then we introduce a new algebraic variant of LWE, Middle-Product Learning With Errors (MP-LWE). We show that this problem is at least as hard as PLWE for many defining polynomials f. As a consequence, any cryptographic system based on MP-LWE remains secure as long as one of these PLWE instances remains hard to solve.

Finally, we illustrate the cryptographic relevance of MP-LWE by building a public-key encryption scheme and a digital signature scheme that are proved secure under the MP-LWE hardness assumption.

ACKNOWLEDGMENTS

Firstly, I want to thank my advisor Damien Stehlé for his excellent mentorship. Thank you for making me discover so many nice ideas, for teaching me not to give up when I was struggling with research problems, for your disponibility to answer all my questions, for encouraging me when I was nervous and for your patience and guidance when I was making mistakes. I am extremely fortunate to have had such an amazing advisor.

I would like to thank Bogdan Dumitru for trusting me when he hired me at Bitdefender, for encouraging me to apply for a PhD and for his constant support throughout these years.

I would like to express my deepest gratitude to my two reviewers Vadim Lyubashevsky and Frederik Vercauteren, who kindly accepted to read my thesis manuscript and gave me many valuable comments and suggestions to improve it. I am also thankful to Philippe Gaborit, Adeline Roux-Langlois and Brigitte Vallée for also accepting to be in my jury. I do feel very lucky to have all of them in my committee.

I thank my coauthors Shi Bai, Dipayan Das, Ryo Hiromasa, Amin Sakzad, Damien Stehlé, Ron Steinfeld, Alexandre Wallet and Zhenfei Zhang for fruitful discussions. I particularly thank Ron and Amin for many things I learned from them throughout these years and for their patience to explain to me a lot of things. Moreover, thank you for inviting me to Melbourne to work with you during my PhD.

I would like to thank my colleagues from the research team of Bitdefender, for the really nice working atmosphere that we have. I particularly want to thank Radu and Mădălina: I really enjoy working with them and I am really fortunate to call them friends. I also want to thank Elena for her constant support and Radu C. for sharing with us his passion for CTF contests. I am also happy to have worked with Dragoş some time ago.

I would also like to thank my colleagues from the AriC team in Lyon for making me feel comfortable during my stay there and from whom I learned a lot of stuff during the seminars: Benoît, Fabien, Alain, Gilles, Bruno, Valentina, Nicolas L., Nicolas B., Weiqiang, Gottfried, Alonso, Hervé, Chitchanok, Elena, Changmin, Fabrice, Huyen, Nathalie, Jean-Michel, Serge, Joris, Claude-Pierre, Alice, Amit, Junqing, Alexandre, Dingding, Laurent, Ida, Guillaume, Vincent, Florent, Anastasia and Octavie. I particularly want to thank Alice for her calm when I was pretending to speak French and for some of the best jokes I have ever heard (some of them about me), Ida for taking me climbing and for our tea breaks, Weiqiang for keeping me updated with crypto gossips, Alexandre for his sarcasm and sense of humor, Gilles for his availability to discuss with me when I was stuck with a research problem, Bruno for giving me constructive feedback after my talks, Valentina for introducing me to French wine and cheese, Huyen, Changmin and Junqing for taking us to some amazing traditional restaurants, Nicolas L. for his effort to teach me French and Nicolas B. for his effort to speak Romanian. I would also like to thank Marie, Chiraz, Nelly, Kadiatou and Myriam for helping me with many administrative issues.

For inviting me to give talks about my work, I would like to thank Adeline Roux-Langlois and Pierre-Alain Fouque in Rennes, Abderrahmane Nitaj and Brigitte Vallée in Caen, Zvika Brakerski, Vinod Vaikuntanathan and Hoeteck Wee in Bertinoro. It was a great opportunity for me to discuss with them and to discover new research environments. I would also like to thank my high school teachers and my university professors for making me discover a lot of nice ideas when I was a student. I especially thank my bachelor and master thesis advisor Marius Vlădoiu for his inspiring math enthusiasm.

Finally, I would like to thank my close family and friends for standing by, my partner for being so supportive while I was working on my thesis and my parents for their love that I will forever feel.

CONTENTS

R	ésum	é		1
\mathbf{A}	bstra	ct		2
A	cknov	wledgn	nents	3
C	onter	\mathbf{nts}		5
Li	st of	\mathbf{symb}	ols	8
R	ésum	é long	en français	9
1	Intr	oducti	ion	15
	1.1	Contri	butions	17
	1.2	Impac	t	19
2	Pre	limina	ries	20
	2.1	Lattic	es	21
		2.1.1	Definitions	21
		2.1.2	Lattice problems	22
		2.1.3	Learning With Errors	23
	2.2	Polyne	omials and structured matrices	24
	2.3	Proba	bilities	27
		2.3.1	Basic definitions	27
		2.3.2	Leftover hash lemma	27
		2.3.3	Gaussian distributions	27
		2.3.4	Gaussian distributions over lattices	28
	2.4	Algeb	raic number theory	29
		2.4.1	Number fields	29
		2.4.2	Embeddings	29
		2.4.3	Rings and ideals in number fields	30
		2.4.4	Ideal lattices	31
		2.4.5	Gaussians on H	31
	2.5	Crypt	ographic definitions	31
		2.5.1	Proofs by reduction	31
		2.5.2	(Quantum) Random-oracle model	31
		2.5.3	Pseudorandom functions	32
		2.5.4	Public-key encryption	32
		2.5.5	Digital signatures	34

3	On	he RLWE and PLWE problems	39						
	3.1	Introduction	. 40						
	3.2	Contributions	. 41						
		3.2.1 Techniques	. 41						
		3.2.2 Related works	. 42						
		3.2.3 Impact	. 43						
		3.2.4 Follow-up work	. 44						
	3.3	Orders in number fields	. 44						
	3.4	3.4 Formal definitions of dual-RLWE, primal-RLWE and PLWE							
	3.5	3.5 From dual-RLWE to primal-RLWE							
	3.6	Controlling the noise growth in the dual to primal reduction	. 47						
	3.7	From primal-RLWE to PLWE	. 49						
		3.7.1 Reducing primal-RLWE to $PLWE^{\sigma}$. 50						
		3.7.2 Distortion between embeddings	. 51						
		3.7.3 A family of polynomials with easily computable distortion	. 51						
		3.7.4 Other "good" families of polynomials	. 54						
	3.8	On small elements and $f'(\alpha)$. 55						
	3.9	Search to decision dual-RIWE	56						
	0.0	3.9.1 A ring-based Leftover Hash Lemma	56						
		3.9.2 Search RIWE to decision RIWE	59						
	3 10	On Vandermonde matrices and the expansion factor	60						
	0.10		. 00						
4	The	Middle-Product Learning With Errors problem	63						
	4.1	Introduction	. 64						
	4.2	Contributions	. 65						
		4.2.1 Follow-up works	. 65						
	4.3	The middle-product of two polynomials	. 66						
	4.4	Middle-product learning with errors	. 67						
	4.5	68 Hardness of MP-LWE							
	4.6	A different way to see the PLWE to MP-LWE reduction	. 70						
	4.7	Hardness of MP-LWE with small secrets	. 72						
-	A		70						
9	App 5 1	Incations of MP-LWE in cryptography	70						
	0.1 5 0	Contributions	· //						
	0.2	Contributions	. 11						
		5.2.1 Related WORKS	. (8						
	٢٩	3.2.2 Follow-up works	. 19						
	0.3	A public-key encryption scheme from MP-LWE	. 80						
		5.3.1 The scheme	. 80						
		5.3.2 Correctness	. 80						
		5.3.3 Security	. 81						
	F 4	5.3.4 Parameters	. 83						
	5.4 5 -	4 An impossibility result for a dual-Kegev scheme based on MP-LWE							
	5.5	A signature scheme based on small secrets MP-LWE	. 84						
		5.5.1 The identification scheme	. 84						
		5.5.2 The signature scheme	. 88						
		5.5.3 Concrete parameters	. 89						
		5.5.4 Implementation \ldots	. 91						
		5.5.5 An attack on Inhomogeneous $PSIS^{\emptyset}$ with small secrets $\ldots \ldots \ldots \ldots \ldots \ldots$. 91						

6 Open problems	94
List of publications	96
Bibligraphy	105
List of figures	107
List of tables	108

LIST OF SYMBOLS

$\mathbb{N},\mathbb{Z},\mathbb{Q},\mathbb{R},\mathbb{C}$	 sets of non-negative integers, integers, rationals, reals,
	complex numbers
$\mathbb{N}^*,\mathbb{Z}^*,\mathbb{Q}^*,\mathbb{R}^*,\mathbb{C}^*$	 sets of non-zero non-negative integers, integers, rationals, reals,
	complex numbers
\mathbb{Z}_q	 $\mathbb{Z}/q\mathbb{Z}$
\mathbb{R}_q	 $\mathbb{R}/q\mathbb{Z}$
x	 column vector
\mathbf{x}^t	 row vector
S^n	 the set of <i>n</i> -dimensional vectors with coefficients in S
$ \mathbf{x} _p$	 ℓ_p norm of x : $(\sum_{i=1}^n x_i ^p)^{1/p}$
$ \mathbf{x} _{\infty}$	 $\ell_{\infty} \text{ norm of } \mathbf{x}: \max_{i} x_{i} $
$\sigma_1(M) \geq \sigma_2(M) \geq \cdots \geq \sigma_n(M)$	 the singular values of the matrix $M \in \mathbb{R}^{n \times n}$
M	 the largest singular value $\sigma_1(M)$ of the matrix M
M _F	 the Frobenius norm of M
$M_{i,j}$	 the element on the <i>i</i> -th row and <i>j</i> -th column from M
[n]	 the set $\{1, 2,, n\}$
arphi(n)	 Euler's totient function evaluated in n
$\log(x)$	 logarithm in base 2 of x
$\lceil x \rceil$	 least integer greater than or equal to x
$x \hookleftarrow D$	 x is sampled from a distribution D
$x \leftrightarrow S \text{ or } x \leftrightarrow U(S)$	 x is sampled from the uniform distribution on the set ${\sf S}$
a b	 the string obtained by the concatenation of the strings \boldsymbol{a} and \boldsymbol{b}
(a,b)	 the greatest common divisor of a and b
$\Sigma \succ 0$	 the set of positive definite matrices
$\Sigma \succ 0$	 Σ is a positive definite matrix
$f(n) = \widetilde{O}(g(n))$	 $f(n) = O(f(n)\log^k(n))$ for some $k > 0$

We use the standard Bachmann–Landau notations and we say that a function ϵ is *negligible* in n if $\epsilon(n) = \frac{1}{n^{\omega(1)}}$.

Résumé long en français

Jusqu'à très récemment, la sécurité des schémas cryptographiques à clé publique reposait en grande partie sur la difficulté présumée de certaines hypothèses de théorie de nombres telles que la factorisation d'entiers [RSA78] ou le calcul du logarithme discret dans des groupes cycliques [DH76]. Shor [Sho94] a fourni en 1994 un algorithme qui permet de résoudre efficacement les deux problèmes ci-dessus avec un ordinateur quantique. Par conséquent, la cryptographie qu'on utilise aujourd'hui pour sécuriser les services de l'Internet ne sera plus sécurisée une fois que des ordinateurs quantiques à grande échelle deviendront pratiques. Le but de la *cryptographie post-quantique* est de développer des schémas cryptographiques qui restent securisés même en présence d'ordinateurs quantiques.

À la fin de l'année 2016, l'Institut National de Standards et Technologie (NIST) a lancé un "processus pour solliciter, évaluer et standardiser" des schémas de chiffrement à clé publique et des signatures digitales qui puissent remplacer les normes actuelles dans un contexte post-quantique. La plupart des propositions fondent leur sécurité sur la difficulté présumée de problèmes bien étudiés portant sur les réseaux, codes, systèmes multivariés ou fonctions de hachage. À l'heure actuelle, parmi les propositions acceptées, trois sur les neuf signatures et neuf sur les dix-sept schémas de chiffrement reposent sur des réseaux.

Réseaux euclidiens

Un réseau est un sous-ensemble de \mathbb{R}^m qui peut être décrit comme l'ensemble de toutes les combinaisons linéaires entières de *n* vecteurs linéairement indépendants $\mathbf{b}_1, \cdots, \mathbf{b}_n$. Nous appelons *n* la dimension du réseau. Le problème le plus connu lié aux réseaux est le problème du plus court vecteur (SVP pour shortest vector problem en anglais) et il demande, étant donné un réseau L, de trouver un vecteur non nul le plus court (pour la norme euclidienne) de L. Le problème bénéficie d'une variante ApproxSVP_{γ} paramétrée par $\gamma > 1$, qui demande de trouver un vecteur non nul dont la norme n'est pas plus que $\gamma \cdot \lambda_1(\mathsf{L})$, où $\lambda_1(\mathsf{L})$ est la norme d'un vecteur non nul le plus court du réseau. Les meilleurs algorithmes connus pour résoudre ApproxSVP_{γ}, l'algorithme de Schnorr [Sch87] et sa version heuristique [SE94], s'exécutent en temps exponentiel dans la dimension *n* du réseau pour des facteurs d'approximation polynomiaux, ce qui rend ApproxSVP_{γ} avec $\gamma = \mathsf{poly}(n)$ approprié comme fondation de sécurité. Pourtant, du point de vue de la conception cryptographique, les problèmes reposant sur des réseaux standards ne sont pas avantageux. La raison est que pour utiliser ces problèmes, il faut choisir un réseau L en particulier, pour lequel le problème correspondant pourrait être facile. Nous appelons souvent ces problèmes *pire-cas*, car ils ne sont pas nécessairement difficiles à résoudre pour tous les réseaux, mais difficiles à résoudre dans le pire des cas.

Cryptographie basée sur les réseaux

La cryptographie basée sur les réseaux repose principalement sur des problèmes *moyen-cas* de réseaux (c'est-à-dire des problèmes définis sur des réseaux pour lesquels des instances aléatoires sont difficiles à résoudre). Dans son travail [Ajt96], Ajtai a montré une connexion remarquable entre les problèmes de

réseaux moyen-cas et pire-cas. Tout d'abord, il a introduit le problème Short Integer Solutions (SIS) qui demande, étant donné certains vecteurs aléatoires $\mathbf{a}_i \in \mathbb{Z}_q^n$, à trouver une combinaison non-triviale d'eux avec des coefficients "courts" dont la somme est zéro. Ajtai a prouvé que pour un certain paramètre γ qui dépend de paramètres du problème SIS, si on peut résoudre SIS en moyenne, il y a un algorithme efficace pour résoudre le problème ApproxSIVP $_{\gamma}$ dans le pire-cas. Pour $\gamma \geq 1$, le problème ApproxSIVP $_{\gamma}$ demande, étant donné un réseau arbitraire n dimensional L, de trouver n vecteurs $\mathbf{v}_1, \ldots, \mathbf{v}_n$ linéairement indépendants et plus courts que $\gamma \cdot \lambda_n(\mathsf{L})$, où $\lambda_n(\mathsf{L})$ est le plus petit réel tel qu'il existe n vecteurs linéairement indépendants dans L de normes inférieures à lui. La difficulté de SIS a été affinée dans une série d'ouvrages ([MR04, GPV08, MP13], etc.) et utilisée comme garantie de sécurité pour la construction de nombreuses applications cryptographiques telles que des fonctions à sens unique et des fonctions résistantes aux collisions ([Ajt96, GGH96, LM06], etc.), des protocoles d'identification et des schémas de signature digitale ([Lyu08, Lyu12], entre autres), etc.

Regev a introduit [Reg05] le problème Learning With Errors (LWE), un problème moyen-cas reposant sur des réseaux qui, contrairement au SIS, convient mieux à la construction de schémas de chiffrement. Le problème LWE est paramétré par des entiers positifs n, q et une distribution d'erreur χ sur \mathbb{R} . La variante de recherche de LWE demande de trouver un secret $\mathbf{s} \in \mathbb{Z}_q^n$ étant donné de nombreux échantillons ($\mathbf{a}_i, b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i \mod q$), où \mathbf{a}_i est choisi uniformément dans \mathbb{Z}_q^n et e_i est tiré selon la distribution χ . Pour un certain γ qui dépend de paramètres du problème, LWE est au moins aussi difficile que les problèmes de réseaux standards tels que ApproxSIVP_{γ} ([Reg05, Pei09, BLP+13], etc.) Le problème LWE a également une version décisionelle dans laquelle on demande de distinguer entre des échantillons du type décrit ci-dessus et des échantillons uniformes sur $\mathbb{Z}_{q}^{n} \times \mathbb{R}/q\mathbb{Z}$. Les deux versions du problème peuvent être réduites l'une à l'autre ([Reg05, Pei09, ACPS09, MP12, BLP+13], etc.). La variante de décision de LWE, qui convient mieux à la conception cryptographique que la variante de recherche LWE, a été initialement utilisée pour créer des schémas de chiffrement à clé publique ([Reg05, GPV08, LP11], etc.). Plus tard, LWE a prouvé sa polyvalence et a également été utilisé pour créer des primitives avancées telles que le chiffrement fonctionnel ([ABCP15, ALS16], etc.), le chiffrement homomorphe ([BV11, BGV12], etc.), le chiffrement reposant sur l'identité ([GPV08, ABB10], etc.), entre autres.

Les schémas cryptographiques basés sur SIS et LWE sont souvent moins efficaces que les protocoles classiques basés sur des hypothèses de théorie des nombres, car ils nécessitent le stockage d'une grande matrice A (qui correspond aux vecteurs \mathbf{a}_i) et le calcul de multiplications matrice-vecteur. *Polynomial Short Integer Solutions* (PSIS)/*Module Short Integer Solutions* (Module-SIS) et *Polynomial Learning With Erors* (PLWE)/*Ring Learning With Erors* (RLWE)/*Module Learning With Erors* (Module-LWE) sont des variantes de SIS, respectivement LWE, qui utilisent des structures algébriques supplémentaires afin d'atteindre une efficacité pratique.

Les problèmes PSIS et PLWE sont paramétrés par un polynôme unitaire irréductible $f \in \mathbb{Z}[x]$ dont le degré est n. Dans le problème PSIS [LM06, PR06] sur l'anneau $\mathbb{Z}_q[x]/(f)$, on reçoit k polynômes choisis uniformément au hasard $\mathbf{a}_1, \ldots, \mathbf{a}_k$ et on doit trouver k éléments $\mathbf{z}_1, \ldots, \mathbf{z}_k$ qui ne soient pas tous nuls et ayant de petits coefficients tels que $\sum_i \mathbf{a}_i \mathbf{z}_i = 0$ dans l'anneau $\mathbb{Z}_q[x]/(f)$. D'autre part, la (variante de décision de) PLWE [SSTX09] demande de distinguer avec une probabilité non-négligeable sur le choix de \mathbf{s} , entre des échantillons de la distribution uniforme sur $\mathbb{Z}_q[x]/(f) \times (\mathbb{R}/q\mathbb{Z})[x]/(f)$ et des échantillons qui dépendent du $\mathbf{s} \in \mathbb{Z}_q[x]/(f)$ choisi uniformément et de certains termes d'erreur avec de petits coefficients. Les deux problèmes peuvent être obtenus comme des variantes spéciales de SIS, respectivement LWE, en identifiant un vecteur $\mathbf{a} \in \mathbb{Z}_q^n$ avec un polynôme dans $\mathbb{Z}_q[x]/(f)$. De plus, cette correspondance permet d'identifier un produit matrice-vecteur par une multiplication de deux polynômes modulo f et d'effectuer une arithmétique rapide, mais également de gagner de l'espace lors du stockage de la matrice \mathbf{A} .

La version de recherche/décision du problème PLWE pourrait être interprétée comme une variante

non-homogène du problème de recherche/décision NTRU. Le problème NTRU est inspiré du système de chiffrement NTRUEncrypt [HPS98], qui est le premier schéma de chiffrement à clé publique reposant sur les réseaux qui utilise des anneaux de polynômes. Il n'y a aucune réduction connue des problèmes pire-cas de réseaux au problème NTRU, ni du problème NTRU à la cryptanalyse de NTRUEncrypt. Pourtant, il existe dans la litérature une version légèrement modifiée [SS11] de NTRUEncrypt qui est prouvée sûre sous l'hypothèse PLWE.

Le problème de recherche/décision RLWE [LPR10] est paramétré par un corps de nombres K et il a deux variantes : primal-RLWE, qui est défini en fonction de l'anneau d'entiers \mathcal{O}_K , et dual-RLWE, qui utilise le dual \mathcal{O}_K^{\vee} . Les problèmes PLWE, primal-RLWE et dual-RLWE sont en fait identiques dans le cas de cyclotomiques (*i.e.* polynômes de degré $\varphi(n)$ dont les seules racines sont les racines primitives *n*-ièmes de l'unité) dont l'ordre *n* est une puissance de 2. Les problèmes PSIS, PLWE et RLWE bénéficient également de réductions des problèmes pire-cas de réseaux [LM06, PR06, SSTX09, LPR10, PRSD17] tels que ApproxSIVP_{γ}, mais dans leur cas, les réseaux impliqués appartiennent à une classe particulière. Nous appelons les réseaux respectifs *réseaux idéaux*, parce qu'un tel réseau correspond à un idéal dans un anneau qui depend du problème.

Les problèmes SIS et PSIS et les problèmes LWE et RLWE peuvent être obtenus comme des cas particuliers de Module-SIS et Module-LWE ([BGV12, LS15]). Les problèmes Module-SIS et Module-LWE sont des problèmes de cas moyen qui sont [LS15] au moins aussi difficiles que le problème ApproxSIVP_{γ} restreint aux réseaux qui correspondent à des modules sur un anneau (*i.e. réseaux modules*).

Les problèmes SIS/LWE sont en fait équivalents à ApproxSIVP_{γ} et les problèmes Module-SIS/LWE sont équivalents à ApproxSIVP_{γ} pour des réseaux modules [AD17]. En revanche, ApproxSIVP_{γ} sur des réseaux idéaux (qui est en fait équivalent à ApproxSVP_{γ} dans ce cas) pourrait être strictement plus facile que PSIS/PLWE/RLWE. Des vulnérabilités potentielles dans la difficulté de ApproxSVP_{γ} pour les réseaux idéaux rendraient les réductions à PSIS/PLWE/RLWE vides de sens. En effet, pour certains polynômes de paramétrage f, tels que les cyclotomiques, le problème ApproxSVP_{γ} pour les réseaux idéaux ([CDPR16, CDW17, PHS19], etc.) est plus facile à résoudre que le problème ApproxSVP_{γ} pour les réseaux généraux, classiquement et quantiquement.

Contributions

Dans cette thèse, nous nous intéressons aux liens entre les variantes algébriques de LWE. Tout d'abord, nous établissons des connexions entre PLWE et RLWE. Deuxièmement, nous introduisons une nouvelle variante algébrique, *Middle-Product Learning With Errors* (MP-LWE), et nous analysons sa relation avec PLWE. Enfin, nous montrons la pertinence cryptographique de MP-LWE.

PLWE et RLWE

Alors que la difficulté du problème de décision RLWE repose sur la difficulté du problème ApproxSVP_{γ} dans les réseaux idéaux de \mathcal{O}_K pour tout corps de nombres K ([PRSD17]), la version décisionnelle de PLWE était connue pour être au moins aussi difficile que ApproxSVP_{γ} dans des réseaux idéaux de $\mathbb{Z}[x]/(f)$ [SSTX09, LPR10] uniquement pour le cas de cyclotomiques d'ordre une puissance de 2. Dans le Chapitre 3, nous montrons que pour de nombreux polynômes f de degré n, les problèmes PLWE, primal-RLWE, dual-RLWE, à la fois pour leurs versions de recherche et de décision, se réduisent (non-uniformément) les uns aux autres en temps polynomial avec des amplifications d'erreurs limitées. En conséquence, la difficulté du problème de décision PLWE n'est plus restreinte à la classe de cyclotomiques d'ordre une puissance de 2, mais plutôt à une classe beaucoup plus large de polynômes de définition.

Nos contributions du Chapitre 3 peuvent être décrites comme suit. Tout d'abord, nous montrons que la réduction de dual-RLWE à primal-RLWE de [LPR10] peut être implémentée avec une petite amplification d'erreur. Cette réduction nécessite la connaissance d'un élément court t dans l'idéal différent $(\mathcal{O}_K^{\vee})^{-1}$ et nous montrons qu'on peut trouver efficacement un tel t dans n'importe quel corps de nombres K par échantillonnage Gaussien. Ensuite, nous étendons ce résultat à une réduction de primal-RLWE à PLWE, mais l'analyse est plus compliquée. Tout d'abord, nous devons gérer la transformation de \mathcal{O}_K à $\mathbb{Z}[x]/(f)$ à l'aide du *conducteur* de $\mathbb{Z}[x]/(f)$. Deuxièmement, nous devons montrer que la réduction n'augmente pas trop l'erreur. Nous décrivons une large classe de polynômes pour lesquels l'augmentation de l'erreur peut être contrôlée. Ces deux réductions sont non-uniformes, car leurs implémentations nécessitent la connaissance d'informations spécifiques sur le corps K. Enfin, nous obtenons une réduction de la variante recherche de RLWE à sa variante décision qui fonctionne pour n'importe quel corps de nombres K, en utilisant la technique dite Oracle Hidden Center Problem introduite en [PRSD17].

Middle-Product Learning With Errors

Il pourrait arriver que PSIS/PLWE/RLWE soient faciles à résoudre pour certains polynômes f (ou corps de nombres K), et difficiles pour les autres. Motivé par cette observation, Lyubashevsky a introduit [Lyu16] une variante de PSIS sur $\mathbb{Z}_q[x]$ et a prouvé que ce nouveau problème (que nous allons appeler PSIS^{\emptyset}) est au moins aussi difficile que le problème PSIS pour tout polynôme de paramétrage f dans une grande famille \mathcal{F} . Par conséquent, tout schéma cryptographique dont la preuve de sécurité repose sur PSIS^{\emptyset} reste sécurisé tant que la recherche de vecteurs courts dans les idéaux de $\mathbb{Z}[x]/(f)$ reste difficile pour au moins un $f \in \mathcal{F}$.

Dans le Chapitre 4, nous définissons un analogue du problème $\mathsf{PSIS}^{\emptyset}$ adapté au contexte de LWE. Nous introduisons *Middle-Product Learning With Errors* (MP-LWE) et montrons que ce problème est au moins aussi difficile que PLWE pour de nombreux polynômes f. Le paramètre d'erreur du problème MP-LWE paramétré par n peut être fixé pour gérer une classe exponentiellement grande de polynômes f. De plus, nous montrons que MP-LWE reste difficile même si les secrets sont tirés d'une distribution qui produit des éléments de petites normes avec une forte probabilité, en réduisant directement le problème PLWE avec des secrets "courts" à MP-LWE avec des secrets "courts". Ce résultat utilise la même technique que le précédent, mais l'analyse des distributions résultantes de secrets et d'erreurs est plus complexe.

Applications de MP-LWE

Dans le Chapitre 5, nous construisons deux primitives cryptographiques dont les preuves de sécurité sont basées sur la difficulté conjecturée du problème MP-LWE. Tout d'abord, nous construisons un schéma de chiffrement à clé publique IND-CPA qui s'inspire du schéma de Regev [Reg09] basé sur LWE. On dit qu'un schéma de chiffrement à clé publique est IND-CPA (pour *indistinguishable against chosen-plaintext attacks* en anglais) si aucun adversaire efficace ne peut reconnaître à quel message clair parmi deux correspond un message chiffré, même si les deux messages ont été choisis par lui-même.

Ensuite, nous construisons un schéma de signature digitale prouvé sûr dans le QROM (pour *quantum* oracle-model en anglais) reposant sur la difficulté conjecturée de MP-LWE avec des secrets "courts". Nous montrons que la signature est UF-CMA (*unforgeable against chosen-message attacks* en anglais), ce qui signifie qu'aucun attaquant, après avoir vu une signature pour n'importe quels messages choisis de manière adaptative, n'est capable de produire une signature valide pour un nouveau message. Nous montrons que pour des paramètres qui atteignent une sécurité similaire à celle utilisée pour instancier la signature de Lyubashevsky [Lyu16], notre signature digitale a des signatures plus courtes d'environ un facteur de 2. Nous prouvons également que la taille de la signature dans [Lyu16] ne peut pas être trop réduite tout en préservant la sécurité du schéma. Par rapport aux schémas de signature à base de réseaux proposés dans le cadre du processus de standardisation du NIST, notre schéma de signature

réalise un compromis risque-performance entre les systèmes à anneau fixe et les systèmes reposant sur LWE.

Nos contributions concernant les relations entre la difficulté de PLWE, RLWE et MP-LWE sont résumées dans la Figure 1. Les contributions sur les relations entre PLWE et RLWE correspondent à la publication [**RSW18**], tandis que les résultats sur MP-LWE ont été publiés dans [**RSSS17**] et [**BDH+20**].

[RSSS17]: Miruna Rosca, Amin Sakzad, Damien Stehlé and Ron Steinfeld. Middle-Product Learning With Errors. In Proc. of CRYPTO, pages 283-297, Springer, 2017.

[RSW18]: Miruna Rosca, Damien Stehlé and Alexandre Wallet. On the Ring-LWE and Polynomial-LWE Problems. In Proc. of EUROCRYPT, pages 146-173, Springer, 2018.

[BDH+20]: Shi Bai, Dipayan Das, Ryo Hiromasa, Miruna Rosca, Amin Sakzad, Damien Stehlé, Ron Steinfeld and Zhenfei Zhang. MPSign: A Signature from Small-Secret Middle-Product Learning with Errors. In Proc. of PKC, pages 66-93, Springer, 2020.

Impact

Les contributions présentées dans cette thèse ont inspiré d'autres travaux. Nous n'en mentionnons maintenant que quelques-uns. Nos résultats présentés au Chapitre 3 ont inspiré une preuve de difficulté alternative de PLWE dans [BBPS19] qui relie le problème à une classe de réseaux différente, correspondant aux idéaux inversibles de l'ordre $\mathbb{Z}[x]/(f)$. Bai *et al.* définissent dans [BBD⁺19] une variante de MP-LWE qui évite la procédure d'échantillonnage Gaussien, le problème Middle-Product Computational Learning With Rounding, et l'utilisent pour construire un schéma de chiffrement à clé publique avec la même efficacité asymptotique que celle du Chapitre 5. Le schéma de chiffrement à clé publique que nous construisons dans le Chapitre 5 a en fait été implémenté et affiné dans [SSZ17, SSZ19] et soumis plus tard au processus de standardisation du NIST. Steinfeld et al. [SSZ17, SSZ19] spécialisent également le résultat de difficulté portant sur MP-LWE à une famille de polynômes qui permet la préservation de la distribution d'erreur dans la réduction du PLWE à MP-LWE afin d'intégrer cette réduction dans la procédure de sélection des paramètres. Lombardi *et al.* [LVV19] proposent un nouveau lemme des restes pour les polynômes sur $\mathbb{Z}_q[x]$ qui ne sont pas réduits modulo un polynôme f et l'utilisent pour construire un schéma de chiffrement reposant sur l'identité dont la preuve de sécurité est basée sur une version légèrement modifiée de MP-LWE. Un cadre général pour analyser toutes les variantes algébriques existantes de LWE, y compris MP-LWE, a été proposé dans [PP19].



Fig. 1: Réductions entre les variantes algébriques de LWE. Chaque flèche peut masquer une dégradation du taux d'erreur (et transfert entre rang du module et module dans le cas de [AD17]). Les flèches noires sans références correspondent aux réductions triviales. Les flèches en pointillés correspondent aux résultats présentés dans le Chapitre 3 et les flèches en tirets correspondent aux résultats présentés dans le Chapitre 4. Les flèches verticales en pointillés correspondent à des réductions non-uniformes. Les réductions impliquant PLWE sont analysées pour une famille restreinte de polynômes de définition.

CHAPTER 1

INTRODUCTION

Until recently, the security of public-key cryptographic schemes was mainly relying on the presumed hardness of some number theoretic assumptions such as factoring [RSA78] or computing discrete logarithms in cyclic groups [DH76]. Shor [Sho94] gave in 1994 an algorithm which allows the solving of the above two problems very fast with a quantum computer. Consequently, the cryptography that we use nowadays on the Internet will be insecure once large scale quantum computers become practical. The goal of *post-quantum cryptography* is to develop cryptographic schemes which remain secure even in the presence of quantum computers.

At the end of 2016, the National Institute of Standards and Technology (NIST) has initiated [NIS] a "process to solicit, evaluate, and standardize" public-key encryption schemes and digital signatures which could replace the current standards in a post-quantum era. Most of the proposals base their security on the conjectured hardness of well-studied problems on lattices, codes, multivariate systems or hash functions. At the moment, among the remaining candidates, 3 out of 9 signatures and 9 out of 17 key encapsulation mechanisms are based on lattices.

Lattices

A lattice is a subset of \mathbb{R}^m which can be described as the set of all integer linear combinations of some linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n$. We call *n* the *dimension* of the lattice. The most famous problem related to lattices is the Shortest Vector Problem (SVP) and it asks, given a lattice L, to find a shortest non-zero vector (with respect to the Euclidean norm) in L. The problem enjoys a relaxed variant ApproxSVP_{γ}, parameterized by $\gamma > 1$, which asks to find a non-zero vector whose norm is no more than $\gamma \cdot \lambda_1(\mathsf{L})$, where $\lambda_1(\mathsf{L})$ is the norm of a shortest non-zero vector in the lattice. The best known algorithms for solving ApproxSVP_{γ}, Schnorr's algorithm [Sch87] and its heuristic version [SE94], run in exponential time in the dimension *n* of the lattice for polynomial approximation factors γ , which makes ApproxSVP_{γ} with $\gamma = \mathsf{poly}(n)$ suitable to be used as security foundation. Still, from a cryptographic design perspective, standard lattice problems are not attractive. The reason is that relying on such problems would mean to pick a particular lattice L, for which the corresponding problem could be easy. We often call such problems *worst-case* problems, because they are not neccessarily hard to solve for any lattice, but hard to solve in the worst-case.

Lattice-based cryptography

Lattice-based cryptography mostly relies on *average-case* lattice problems (*i.e.* problems defined over lattices for which random instances are hard to solve). In his seminal work [Ajt96], Ajtai proved an incredible connection between average-case and worst-case lattice problems. He first introduced the Short Integer Solutions (SIS) problem which asks, given some uniformly random vectors $\mathbf{a}_i \in \mathbb{Z}_q^n$, to find a non-trivial "short" integer combination of them which sums to zero. Ajtai proved that for a certain parameter γ depending on the parameters of the SIS problem, solving SIS on average would imply an efficient algorithm for solving the ApproxSIVP_{γ} problem in the worst case. For $\gamma \geq 1$, the ApproxSIVP_{γ} problem asks, given an *n*-dimensional lattice L, to find *n* vectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ linearly independent and shorter than $\gamma \cdot \lambda_n(\mathsf{L})$, where $\lambda_n(\mathsf{L})$ is the smallest number with the property that there exist *n* linearly independent vectors in L of norm less than it. The hardness of SIS has been refined in a series of works ([MR04, GPV08, MP13], etc.) and used as security guarantee for building many cryptographic applications such as one-way and collision-resistant hash functions ([Ajt96, GGH96, LM06], etc.), identification protocols and digital signature schemes ([Lyu08, Lyu12], among others), etc.

Regev introduced in [Reg05] the Learning With Errors (LWE) problem, an average-case lattice problem which is, contrary to SIS, more suitable for building encryption schemes. The LWE problem is parameterized by positive integers n, q, and an error distribution χ over \mathbb{R} . The (*search* variant of the) LWE problem asks to find the secret $\mathbf{s} \in \mathbb{Z}_q$ given many samples ($\mathbf{a}_i, b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i \mod q$), where \mathbf{a}_i is uniformly random in \mathbb{Z}_q^n and e_i is drawn from the distribution χ . For some γ dependent on the parameters of the problem, LWE is at least as hard as standard lattice problems such as ApproxSIVP_{γ} ([Reg05, Pei09, BLP⁺13], etc.). The LWE problem also has a *decision* version which asks to distinguish between samples of the type described above and uniform samples on $\mathbb{Z}_q^n \times \mathbb{R}/q\mathbb{Z}$. The two versions of the problem can be reduced one to the other ([Reg05, Pei09, ACPS09, MP12, BLP⁺13], etc.). The decision variant of LWE, which is more suitable for cryptographic design than search LWE, has been initially used to build public-key encryption schemes ([Reg05, GPV08, LP11], etc.). Later on, LWE proved its versatility and it has been used to also build advanced primitives such as functional encryption ([ABCP15, ALS16], etc.), homomorphic encryption ([BV11, BGV12], etc.), identity-based encryption ([GPV08, ABB10], etc.), among others.

The cryptographic schemes based on SIS and LWE are usually less efficient than the classical protocols based on number theoretic assumptions, because they require the storage of a big matrix A (corresponding to the vectors \mathbf{a}_i) and the computation of several matrix-vector multiplications. Polynomial Short Integer Solutions (PSIS)/Module Short Integer Solutions (Module-SIS) and Polynomial Learning With Errors (PLWE)/Ring Learning With Errors (RLWE)/Module Learning With Errors (Module-LWE) are variants of SIS, respectively LWE, which make use of extra algebraic structures in order to reach practical efficiency.

The PSIS and PLWE problem are both parameterized by a monic irreducible polynomial $f \in \mathbb{Z}[x]$ of degree *n*. In the PSIS problem [LM06, PR06] over the ring $\mathbb{Z}_q[x]/(f)$, one is given *k* uniformly random polynomials $\mathbf{a}_1, \ldots, \mathbf{a}_k$ and is asked to find *k* not all zero elements $\mathbf{z}_1, \ldots, \mathbf{z}_k$ having small coefficients such that $\sum_i \mathbf{a}_i \mathbf{z}_i = 0$ in the ring $\mathbb{Z}_q[x]/(f)$. On the other hand, the (decision variant of the) PLWE problem [SSTX09] asks to distinguish with non-negligible probability over the choice of \mathbf{s} , between arbitrarily many samples from the uniform distribution on $\mathbb{Z}_q[x]/(f) \times (\mathbb{R}/q\mathbb{Z})[x]/(f)$ and samples which depend on the uniformly random $\mathbf{s} \in \mathbb{Z}_q[x]/(f)$ and on some error terms with small coefficients. Both problems can be obtained as special variants of SIS, respectively LWE, by identifying a vector $\mathbf{a} \in \mathbb{Z}_q^n$ with a polynomial in $\mathbb{Z}_q[x]/(f)$. This correspondence allows to further identify a matrix-vector product by a multiplication of two polynomials modulo *f* and perform fast-arithmetic, but to also save space when storing the matrix A.

The search/decision version of the PLWE problem could be interpreted as the inhomogeneous variant of the search/decision NTRU problem. The NTRU problem is inspired by the NTRUEncrypt cryptosystem [HPS98], which is the first lattice-based public-key encryption scheme using polynomial rings. There is no known reduction from worst-case lattice problems to the NTRU problem, nor from the NTRU problem to breaking the security of the respective cryptosystem. Still, there exists in the literature a slightly modified version [SS11] of NTRUEncrypt which is provably secure under the PLWE hardness assumption.

The search/decision RLWE problem [LPR10] is parameterized by a number field K and has two

variants: primal-RLWE, which is defined based on the ring of integers \mathcal{O}_K , and dual-RLWE, which uses the dual \mathcal{O}_K^{\vee} . The PLWE, primal-RLWE and dual-RLWE problems are actually identical in the case of cyclotomics (*i.e.* polynomials of degree $\varphi(n)$ whose only roots are the *n*-th primitive roots of unity) whose order *n* is a power of two. The PSIS, PLWE and RLWE problems also enjoy reductions from worst-case lattice problems [LM06, PR06, SSTX09, LPR10, PRSD17] such as ApproxSIVP_{γ}, but in their case the lattices involved belong to a special class. We call the respective lattices *ideal lattices*, because any such lattice corresponds to an ideal of a ring depending on the problem.

The SIS and PSIS problems, respectively the LWE and RLWE problems, can be obtained as particular cases of Module-SIS and Module-LWE ([BGV12, LS15]). The Module-SIS and Module-LWE problems are average-case problems proven [LS15] to be at least as hard as the ApproxSIVP_{γ} problem on lattices which correspond to modules over a specific ring (*i.e. module lattices*).

The SIS/LWE problems are actually equivalent to $ApproxSIVP_{\gamma}$ and the Module-SIS/LWE problems equivalent to $ApproxSIVP_{\gamma}$ over module lattices [AD17]. On the contrary, $ApproxSIVP_{\gamma}$ over ideal lattices (which is actually equivalent to $ApproxSVP_{\gamma}$ in this case) could be strictly easier than PSIS/PLWE/RLWE. Potential weaknesses in the hardness of $ApproxSVP_{\gamma}$ over ideal lattices would make the reductions to PSIS/PLWE/RLWE vacuous. For some polynomials f such as cyclotomics, the $ApproxSVP_{\gamma}$ problem is indeed easier on ideal lattices ([CDPR16, CDW17, PHS19], etc.) than on general lattices, both classically and quantumly.

1.1 Contributions

In this thesis, we study relationships between algebraic variants of LWE. First, we establish connections between PLWE and RLWE. Second, we introduce a new algebraic variant, the Middle-Product Learning With Errors (MP-LWE) problem, and discuss its relationship with PLWE. Last, we illustrate the cryptographic use of MP-LWE.

PLWE and RLWE

While the hardness of decision RLWE relies now on the hardness of solving ApproxSVP_{γ} in lattices corresponding to \mathcal{O}_K ideals for any field K due to [PRSD17], the decision PLWE problem was known to be at least as hard as ApproxSVP_{γ} in $\mathbb{Z}[x]/(f)$ ideal lattices [SSTX09, LPR10] only for the case of cyclotomics of order a power of 2. In Chapter 3, we show that for exponentially many polynomials fof degree n, the corresponding PLWE, primal-RLWE, dual-RLWE problems, both in their search and decision versions, (non-uniformly) reduce to one another in polynomial time with limited error rate increases. As a consequence, the hardness of decision PLWE is not restricted anymore to the class of cyclotomics of order a power of two, but rather to a much larger class of polynomials.

On a high level, our contributions from Chapter 3 can be described as follows. We first show that the reduction from dual-RLWE to primal-RLWE from [LPR10] can be implemented with a small error growth. This reduction requires the knowledge of a short element t in the different ideal $(\mathcal{O}_K^{\vee})^{-1}$ and we show that we can efficiently find such a t in any field K by Gaussian sampling. Then, we extend this result to a reduction from primal-RLWE to PLWE, but the analysis is more complicated. Firstly, we have to handle the transformation from \mathcal{O}_K to $\mathbb{Z}[x]/(f)$ with the help of the so-called *conductor* of $\mathbb{Z}[x]/(f)$. Secondly, we have to show that the reduction does not increase the error too much. We describe a huge class of polynomials for which the error increase can be controlled. These two reductions are non-uniform, in the sense that their implementations require the knowledge of specific information on the field K. Finally, we obtain a search to decision reduction for RLWE which works for any number field K by using the so-called *Oracle Hidden Center Problem* technique introduced in [PRSD17].

Middle-Product Learning With Errors

It could happen that $\mathsf{PSIS}/\mathsf{PLWE}/\mathsf{RLWE}$ are easy to solve for some polynomials f (or number fields K), and hard for others. Motivated by this, Lyubashevsky introduced in [Lyu16] a variant of PSIS over $\mathbb{Z}_q[x]$ and proved that this new problem (which we are going to refer to as $\mathsf{PSIS}^{\emptyset}$) is at least as hard as the PSIS problem for any parameterizing polynomial f in a large family \mathcal{F} . As a consequence, any cryptographic scheme whose security proof relies on $\mathsf{PSIS}^{\emptyset}$ stays secure as long as finding short vectors in ideals of $\mathbb{Z}[x]/(f)$ remains hard for at least one $f \in \mathcal{F}$.

In Chapter 4, we define an analogue of the PSIS^\emptyset problem adapted to the LWE context. We introduce the Middle-Product Learning With Errors (MP-LWE) problem and show that this problem is at least as hard as PLWE for many polynomials f. The noise parameter of the MP-LWE problem of parameter ncan be set to handle an exponentially large class of polynomials f. We further show that MP-LWE remains hard even if the secrets are drawn from a distribution which produces small elements with high probability, by directly reducing the PLWE problem with "short" secrets to MP-LWE with "short" secrets. This result follows the same blueprint as the previous one, but the analysis of the resulting error and secret distributions is more involved.

Applications of MP-LWE

In Chapter 5, we build two cryptographic primitives whose security proofs are based on the conjectured hardness of the MP-LWE problem. First, we build an IND-CPA public-key encryption scheme which follows the same blueprint as Regev's scheme [Reg09] based on LWE. We say that a public-key encryption scheme is IND-CPA (indistinguishable against chosen-plaintext attacks) if no efficient adversary can recognize which of two messages is encrypted in a given ciphertext, even if the two messages have been chosen by itself.

Then, we build a digital signature scheme tightly secure in the quantum random-oracle model under the conjectured hardness of the MP-LWE with small secrets assumption. We show that the signature is unforgeable against chosen-message attacks (UF-CMA), which means that no attacker, after having seen (possibly more than) one signature for any poly(n) adaptively chosen messages, is able to produce a valid signature for a new message. We show that for parameters that achieve similar security to those used to instantiate Lyubashevsky's signature [Lyu16], our digital signature has shorter signatures by approximately a factor of 2. We also provide evidence that the signature size in [Lyu16] cannot be decreased too much while preserving the security of the scheme. Compared to the latticebased signature schemes proposed for standardization, our signature scheme achieves a risk-performance tradeoff between fixed-ring and LWE-based schemes.

Our contributions regarding the relationships between the hardness of PLWE, RLWE and MP-LWE are summarized in Figure 1.1. The contributions on relationships between PLWE and RLWE correspond to [**RSW18**], while the results on MP-LWE have been published in [**RSSS17**] and [**BDH+20**].

[RSSS17]: Miruna Rosca, Amin Sakzad, Damien Stehlé and Ron Steinfeld. Middle-Product Learning With Errors. In Proc. of CRYPTO, pages 283-297, Springer, 2017.

[**RSW18**]: Miruna Rosca, Damien Stehlé and Alexandre Wallet. On the Ring-LWE and Polynomial-LWE Problems. In Proc. of EUROCRYPT, pages 146-173, Springer, 2018.

[**BDH+20**]: Shi Bai, Dipayan Das, Ryo Hiromasa, Miruna Rosca, Amin Sakzad, Damien Stehlé, Ron Steinfeld and Zhenfei Zhang. MPSign: A Signature from Small-Secret Middle-Product Learning with Errors. In Proc. of PKC, pages 66-93, Springer, 2020.

1.2 Impact

The contributions presented in this thesis inspired other works. We mention now just a few of them. Our results presented in Chapter 3 inspired an alternative hardness proof of PLWE in [BBPS19] by relating the problem to a different class of lattices, corresponding to the invertible ideals of the order $\mathbb{Z}[x]/(f)$. Bai *et al.* introduced in [BBD⁺19] a variant of MP-LWE which avoids the Gaussian sampling procedure, the Middle-Product Computational Learning With Rounding problem, and used it to build a public-key encryption scheme with the same asymptotic efficiency as the one from Chapter 5. The public-key encryption scheme that we build in Chapter 5 has been implemented and refined in [SSZ17, SSZ19] and later submitted to the NIST standardization process. Steinfeld *et al.* [SSZ17, SSZ19] also specialize the hardness result on MP-LWE to a family of polynomials which allow the preservation of the noise distribution in the PLWE to MP-LWE reduction in order to incorporate the reduction into the parameter selection procedure with a limited efficiency loss. Lombardi *et al.* [LVV19] proposed a new leftover-hash lemma for polynomials over $\mathbb{Z}_q[x]$ that are not folded modulo some polynomial f and used it to build an identity-based encryption scheme whose proof of security is based on a slightly modified version of MP-LWE. A general framework to analyze all the existing algebraic variants of LWE, including MP-LWE, has been proposed in [PP19].



Fig. 1.1: Reductions between algebraic variants of LWE. Each arrow may hide a noise rate degradation (and module rank - modulus magnitude transfer in the case of [AD17]). The black arrows without references correspond to trivial reductions. The dotted arrows correspond to the results presented in Chapter 3 and the dashed arrows correspond to the results presented in Chapter 4. The vertical dotted arrows correspond to non-uniform reductions. The reductions involving PLWE are analyzed for a restricted family of defining polynomials.

Chapter 2

Preliminaries

In this chapter, we recall some preliminary results on lattices, polynomials and their relationship with structured matrices, probabilities and algebraic number theory. We then introduce the cryptographic definitions that we are going to use in this thesis.

Contents

2.1	Latt	ices	21
	2.1.1	Definitions	21
	2.1.2	Lattice problems	22
	2.1.3	Learning With Errors	23
2.2	Poly	nomials and structured matrices	24
2.3	Prob	pabilities	27
	2.3.1	Basic definitions	27
	2.3.2	Leftover hash lemma	27
	2.3.3	Gaussian distributions	27
	2.3.4	Gaussian distributions over lattices	28
2.4	Alge	braic number theory	29
	2.4.1	Number fields	29
	2.4.2	Embeddings	29
	2.4.3	Rings and ideals in number fields	30
	2.4.4	Ideal lattices	31
	2.4.5	Gaussians on H	31
2.5	Cryp	otographic definitions	31
	2.5.1	Proofs by reduction	31
	2.5.2	(Quantum) Random-oracle model	31
	2.5.3	Pseudorandom functions	32
	2.5.4	Public-key encryption	32
	2.5.5	Digital signatures	34

2.1 Lattices

In this section, we briefly recall the definitions and results on lattices that we are going to use in the next chapters. There are several resources which can be consulted for a good introduction on lattices and lattice-based cryptography. Among them, we recommend the lecture notes [Pei, Reg] and the two surveys [MR09] and [Pei15].

2.1.1 Definitions

We start by giving the definition of a lattice.

Definition 2.1. Given m linearly independent vectors $(\mathbf{b}_i)_{1 \leq i \leq m}$ over \mathbb{R}^n , the lattice $\mathsf{L} \subseteq \mathbb{R}^n$ generated by them is the set of all linear integer combinations of the vectors \mathbf{b}_i :

$$\mathsf{L}(\mathbf{b}_1,\ldots,\mathbf{b}_m) = \{\sum_{i=1}^m z_i \mathbf{b}_i : z_1,\ldots,z_m \in \mathbb{Z}\}.$$

The set $(\mathbf{b}_i)_{1 \leq i \leq m}$ is called a *basis* of the lattice L. If we define B as the $n \times m$ matrix whose columns are $\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_m$, we simply write $\mathsf{L} = \mathsf{L}(\mathsf{B}) = \mathsf{L}(\mathbf{b}_1, \ldots, \mathbf{b}_m)$. We say that m is the *rank* of the lattice L and n is its *dimension*. If n = m, the lattice is called *full-rank*. In this thesis, all the lattices will be full-rank. A lattice has many bases and any two bases $\mathsf{B}_1, \mathsf{B}_2 \in \mathbb{R}^{n \times m}$ generate the same lattice if and only if $\mathsf{B}_2 = \mathsf{B}_1 \cdot \mathsf{U}$ for some integer matrix U whose determinant is 1 or -1.



Fig. 2.1: A two-dimensional lattice and two of its bases: $\{\mathbf{b}_1, \mathbf{b}_2\}$ and $\{\mathbf{v}_1, \mathbf{v}_2\}$.

Definition 2.2. If L is a lattice and $(\mathbf{b}_i)_{1 \le i \le m}$ is a basis of it, we define the determinant of L as $\det(L) = (\det(<\mathbf{b}_i,\mathbf{b}_j>)_{i,j})^{1/2}$.

Although the determinant is defined by first fixing a basis B of the lattice, it can be shown that its value is actually independent on the choice of B.

For any $p \in \mathbb{N}^* \cup \{\infty\}$ and any lattice $\mathsf{L} \subset \mathbb{R}^n$, $\lambda_1^p(\mathsf{L})$ denotes the ℓ_p norm of a shortest nonzero vector in L . When p = 2 we usually simply write $\lambda_1(\mathsf{L})$.

Theorem 2.1 (Minkowski's First Theorem). For any full-rank lattice $L \in \mathbb{R}^n$, we have $\lambda_1(L) \leq \sqrt{n} \cdot \det(L)^{1/n}$.

Definition 2.3. Given a lattice $L \subset \mathbb{R}^n$, we define its dual lattice L^* as $L^* = \{ \mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}$ for all $\mathbf{y} \in L \}$.

2.1.2 Lattice problems

Minkowski's bound (Theorem 2.1) on the norm of a shortest nonzero vector in a lattice L is in general loose and we may ask whether we can actually compute or better approximate $\lambda_1(L)$ or find a lattice vector whose norm is $\lambda_1(L)$. In the next problem definitions, $\gamma \geq 1$ is a parameter of the problems.

Definition 2.4 (ApproxSVP_{γ}). Given a lattice $L \subset \mathbb{R}^n$, find a nonzero vector $\mathbf{v} \in L$ such that $||\mathbf{v}||_2 \leq \gamma \cdot \lambda_1(L)$.

Definition 2.5 (GapSVP_{γ}). Given a lattice $L \subset \mathbb{R}^n$ and a positive integer d, distinguish between the two cases $\lambda_1(L) \leq d$ and $\lambda_1(L) > \gamma \cdot d$.

For $\gamma = 1$, GapSVP_{γ} and $\mathsf{ApproxSVP}_{\gamma}$ are equivalent (see e.g. [MG02]). For $\gamma > 1$, GapSVP_{γ} trivially reduces to $\mathsf{ApproxSVP}_{\gamma}$ ([MG02, p. 20]), but the converse is not in general true. Still, there exists a dimension preserving randomized reduction [SD16, Che13] from $\mathsf{ApproxSVP}_{\gamma}$ to $\mathsf{GapSVP}_{\gamma'}$ for $\gamma' = \gamma^{O(n/\log n)}$.



Fig. 2.2: The vector \mathbf{v} is a shortest nonzero vector in $\mathbf{L} = \mathbf{L}(\mathbf{v}_1, \mathbf{v}_2)$ and \mathbf{v}' is a shortest vector up to $\gamma = 1.5$ approximation factor.

Another well-known lattice problem is $\mathsf{ApproxCVP}_{\gamma}$. This problem also enjoys a decision variant GapCVP_{γ} , defined analogously to GapSVP_{γ} , and is closely related to $\mathsf{ApproxSVP}_{\gamma}$ in the sense that $\mathsf{ApproxSVP}_{\gamma}$ reduces to $\mathsf{ApproxCVP}_{\gamma}$ ([GMSS99]) for any $\gamma \geq 1$. Moreover, there exists a folklore reduction from $\mathsf{ApproxCVP}_{\sqrt{n}\cdot\gamma^2}$ to $\mathsf{ApproxSVP}_{\gamma}$ whose proof can be found in [SD15].

Definition 2.6 (Approx CVP_{γ}). Given a lattice $\mathsf{L} \subset \mathbb{R}^n$ and a vector $\mathbf{t} \in \mathbb{R}^n$, find a vector $\mathbf{v} \in \mathsf{L}$ such that $||\mathbf{v} - \mathbf{t}||_2 \leq \gamma \cdot \mathsf{dist}(\mathbf{t},\mathsf{L})$, where $\mathsf{dist}(\mathbf{t},\mathsf{L}) := \inf\{||\mathbf{t} - \mathbf{t}'||_2 : \mathbf{t}' \in \mathsf{L}\}$.

Notice that the above problem definitions are given in terms of the ℓ_2 norm, but we could define these problems with respect to any other norm. When $\gamma = 1$, we call the corresponding problems *exact* and usually omit the subscript γ .

The fastest known algorithms to solve the above lattice problems, sieving ([AKS01, ADRSD15, ASD18], etc.) and enumeration ([Kan83], etc.), run in exponential time in the dimension of the lattice for polynomial approximation factors, which makes these computational problems suitable to be used as hardness foundation for cryptographic schemes.



Fig. 2.3: The vector \mathbf{w} is the closest vector to \mathbf{t} belonging to $\mathbf{L} = \mathbf{L}(\mathbf{v}_1, \mathbf{v}_2)$.

2.1.3 Learning With Errors

In [Reg05], Regev introduced a natural generalization of the so-called Learning Parity with Noise problem [BFKL93, BKW03]: the Learning With Errors problem (LWE). He proved its hardness based on the presumed hardness of GapSVP_{γ} and illustrated its cryptographic use by constructing a public-key encryption scheme whose security relies on the hardness of LWE. Following [Reg05], several (advanced) cryptographic schemes have been built using LWE ([BCD⁺16, GPV08], etc.). In this thesis, we study algebraic variants of LWE, but we recall LWE now for the sake of completeness. The LWE problem is parameterized by two integers $n \geq 1$ and $q \geq 2$ and an error distribution χ on \mathbb{R} and relies on the following distribution.

Definition 2.7 (LWE distribution). For a vector $\mathbf{s} \in \mathbb{Z}_q^n$ called the secret, we define the LWE distribution $\mathcal{D}_{\mathbf{s},\chi}$ as the distribution over $\mathbb{Z}_q^n \times \mathbb{R}_q$ obtained by sampling $\mathbf{a} \leftarrow U(\mathbb{Z}_q^n)$, $e \leftarrow \chi$ and returning the pair $(\mathbf{a}, b = < \mathbf{a}, \mathbf{s} > + e \mod q)$.

The LWE problem has two variants: search and decision. The decision variant is more suitable for cryptographic purposes than the search variant.

Definition 2.8 (Search LWE_{q,n, χ}). Given many samples $(\mathbf{a}_i, b_i) \in \mathbb{Z}_q^n \times \mathbb{R}_q$ from the distribution $\mathcal{D}_{\mathbf{s},\chi}$ for some $\mathbf{s} \in \mathbb{Z}_q^n$, find the secret \mathbf{s} .

The search variant of LWE could be seen as an average-case exact ApproxCVP problem on the lattice $\mathcal{L}(A) := \{A \cdot \mathbf{w} : \mathbf{w} \in \mathbb{Z}_q^n\} + q\mathbb{Z}^m$ where every row of the matrix $A \in \mathbb{Z}_q^{m \times n}$ corresponds to a sample \mathbf{a}_i . Indeed, for a typical choice of parameters, the vector \mathbf{b} (whose entries are the b_i 's) is very close to one of the vectors of $\mathcal{L}(A)$ and the goal in search LWE is exactly to recover that vector.

Definition 2.9 (Decision LWE_{q,n, χ}). Decision LWE_{q,n, χ} consists in distinguishing between a sampler from $\mathcal{D}_{\mathbf{s},\chi}$ and a uniform sampler over $\mathbb{Z}_q^n \times \mathbb{R}_q$, with non-negligible probability over the choice of \mathbf{s} .

Under some conditions on the parameters ([Reg05, Pei09, MM11, MP12], etc.), the search and decision variants are actually equivalent. Notice that both the search and decision LWE problems are easy to solve if the error distribution χ always outputs 0, since we can efficiently recover the secret s using Gaussian elimination. We refer to the above search/decision variants of LWE as *continuous* since the error distribution χ is continuous. There also exist corresponding *discrete* variants, where χ is a

distribution on $\mathbb{Z}_q := \mathbb{Z}/q\mathbb{Z}$. The main result in [Reg05] is a quantum reduction from solving certain lattice problems to solving the LWE problem.

Theorem 2.2 ([Reg05]). For any $0 < \alpha < 1$, any $q \ge 2$ and any Gaussian error distribution χ of standard deviation $\alpha q \ge 2\sqrt{n}$, solving search $\mathsf{LWE}_{q,n,\chi}$ is at least as hard as quantumly solving GapSVP_{γ} on arbitrary n-dimensional lattices, for some $\gamma = \widetilde{O}(n/\alpha)$.

We recall now some of the follow-up works where the hardness of LWE has been analyzed in different contexts and parameter settings. In [Pei09], Peikert analyzed the LWE problem in the setup where the error distribution is a continuous Gaussian. First, he showed that the search variant of LWE with exponential modulus remains at least as hard as $GapSVP_{\gamma}$ for $\gamma = \tilde{O}(n/\alpha)$ even classically. Secondly, he gave a classical hardness proof of the search variant of LWE with polynomial modulus based on a non-standard lattice problem. He also gave a search-to-decision LWE reduction which requires the modulus q to be a product of distinct and sufficiently large polynomially bounded primes. Later on, building upon [Pei09], Brakerski *et al.* [BLP⁺13] proved the classical hardness of LWE in dimension n with polynomial modulus based on the hardness of $GapSVP_{\gamma}$ in dimension $\simeq \sqrt{n}$ using the so-called *modulus reduction* technique.

2.2 Polynomials and structured matrices

In this section, we exhibit some connections between polynomials and the structured matrices that we are going to use in this thesis and recall the definition of the expansion factor of a polynomial [LM06]. We refer the reader to [Pan01] for more details on structured matrices.

Let R be a ring. We let R[x] denote the set of polynomials with coefficients in R and for any k > 0, we let $R^{\leq k}[x]$ denote the set of polynomials in R[x] of degree < k. Given a polynomial $f(x) = f_0 + f_1 x + \cdots + f_{k-1} x^{k-1} \in R^{\leq k}[x]$, we call f_0 the free coefficient of f and f_{k-1} the leading coefficient of f. We use the following notations: $\overline{f}(x) = f_{k-1} + f_{k-2}x + \cdots + f_0x^{k-1} \in R^{\leq k}[x]$, $\mathbf{f} = (f_0, \ldots, f_{k-1})^T \in R^k$ and $\overline{\mathbf{f}} = (f_{k-1}, \ldots, f_0)^T \in R^k$. Notice that $\overline{f}(x) = f(\frac{1}{x}) \cdot x^{k-1}$. Also, the definition of \overline{f} takes into consideration the set where the polynomial f lives and not the actual degree of f. For any two integers $0 \leq a \leq b$, we let $[f(x)]_a^b$ denote the polynomial $f_a + f_{a+1}x + \cdots + f_b x^{b-a}$. When a = b, we simply write $[f(x)]_a$ (or f_a) instead of $[f(x)]_a^a$. For any integers $0 < a \leq b$ and any polynomials $p \in R^{\leq a}[x]$ and $q \in R^{\leq b}[x]$, we consider their product $pq \in R^{\leq a+b-1}[x]$ and their sum $p + q \in R^{\leq b}[x]$. We have that $\overline{pq}(x) = \overline{p}(x) \cdot \overline{q}(x)$ and $\overline{p+q}(x) = \overline{p}(x) \cdot x^{b-a} + \overline{q}(x)$.

We let R[[x]] denote the ring of formal power series in x with coefficients in the ring R. We extend the above notation $[f(x)]_a^b$ to any formal series $f(x) \in R[[x]]$.

If R^k is a normed vector space over \mathbb{R} , for any $p \in \mathbb{N}^* \cup \{\infty\}$ we extend the definition of the ℓ_p norm of vectors to polynomials and define $||f(x)||_p := ||\mathbf{f}||_p$.

When the indeterminate x is clear from the context, we simply write f instead of f(x).

Definition 2.10. Let f be a polynomial of degree $m \ge 0$. For any d > 0 and any $a \in R[x]$, we let $\operatorname{Rot}_{f}^{d}(a)$ denote the matrix in $R^{d \times m}$ whose *i*-th row is given by the coefficients of the polynomial $(x^{i-1} \cdot a) \mod f$, for any $i = 1, \ldots, d$. When d = m, we will use the notation $\operatorname{Rot}_{f}(a)$ instead of $\operatorname{Rot}_{f}^{m}(a)$.

For example, the $\operatorname{Rot}_{f}^{m}(a)$ matrix associated to the polynomials $a(x) = a_0 + a_1 x + \ldots + a_{m-1} x^{m-1}$

and $f(x) = x^m + 1$ is:

$$\mathsf{Rot}_{f}(a) = \begin{pmatrix} a_{0} & a_{1} & \dots & a_{n-2} & a_{m-1} \\ -a_{m-1} & a_{0} & \dots & a_{m-3} & a_{m-2} \\ \vdots & & \ddots & & \vdots \\ \vdots & & \ddots & & \vdots \\ -a_{1} & -a_{2} & \dots & -a_{m-1} & a_{0} \end{pmatrix}$$

Note that if $a' = a \mod f$, then $\operatorname{Rot}_f^d(a) = \operatorname{Rot}_f^d(a')$ for any d and $\operatorname{Rot}_f(a \cdot b) = \operatorname{Rot}_f(a) \cdot \operatorname{Rot}_f(b)$ for any $a, b \in R[x]$.

Definition 2.11. Let f be a polynomial of degree m and d > 0. We define \mathbf{M}_{f}^{d} as the (Hankel) matrix in $\mathbb{R}^{d \times m}$ such that for any $1 \leq i \leq d$ and $1 \leq j \leq m$, the coefficient $(\mathbf{M}_{f})_{i,j}$ is the constant coefficient of $x^{i+j-2} \mod f$. When d = m, we simply write \mathbf{M}_{f} instead of \mathbf{M}_{f}^{m} .

For example, the \mathbf{M}_f matrix associated to $f = x^m + 1$ is:

$$\mathbf{M}_{f} = \begin{pmatrix} 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & -1 \\ 0 & 0 & \dots & -1 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & -1 & \dots & 0 & 0 \end{pmatrix}$$

The matrix \mathbf{M}_f helps rewriting multiplication on the left by matrix $\mathsf{Rot}_f(a)$ as a multiplication on the right by **a**.

Lemma 2.1. For any $a \in R^{\leq m}[x]$, we have $\operatorname{Rot}_f(a) \cdot (1, 0, \dots, 0)^T = \mathbf{M}_f \cdot \mathbf{a}$.

Proof. First, the *i*-th coordinate of the left hand side is the constant coefficient of $x^{i-1} \cdot a \mod f$. Second, the *i*-th coordinate of the right hand side is

$$((a_0 x^{i-1} \mod f) \mod x) + \dots + ((a_{m-1} x^{m+i-2} \mod f) \mod x),$$

which can be re-written as $x^{i-1}(a_0 + \cdots + a_{m-1}x^{m-1} \mod f) \mod x = (x^{i-1} \cdot a \mod f) \mod x$. The latter is the constant coefficient of $x^{i-1} \cdot a \mod f$.

Definition 2.12. For any d, k > 0 and $a \in R^{\langle k}[x]$, we let $\mathsf{Toep}^{d,k}(a)$ denote the matrix in $R^{d \times (k+d-1)}$ whose *i*-th row, for i = 1, ..., d, is given by the coefficients of $x^{i-1} \cdot a$:

$$\mathsf{Toep}^{d,k}(a) = \begin{pmatrix} a_0 & a_1 & \dots & a_{k-1} & 0 & \dots & \dots & 0 \\ 0 & a_0 & \dots & a_{k-2} & a_{k-1} & \dots & \dots & 0 \\ \vdots & & \ddots & & & \ddots & & \\ \vdots & & & \ddots & & & \ddots & \\ 0 & 0 & \dots & \dots & a_0 & \dots & a_{k-2} & a_{k-1} \end{pmatrix}$$

Lemma 2.2. For any d, k > 0 and any $a \in \mathbb{R}^{\leq k}[x]$, we have $\mathsf{Rot}_f^d(a) = \mathsf{Toep}^{d,k}(a) \cdot \mathsf{Rot}_f^{k+d-1}(1)$.

Proof. It is sufficient to prove that the rows of $\mathsf{Rot}_f^d(a)$ and $\mathsf{Toep}^{d,k}(a) \cdot \mathsf{Rot}_f^{k+d-1}(1)$ are equal. We just note that the *i*-th row of $\mathsf{Rot}_f^{k+d-1}(1)$ is $x^{i-1} \mod f$, for $i = 1, \ldots, k+d-1$ and these will fill the gap in the definitions of $\mathsf{Rot}_f^d(a)$ and $\mathsf{Toep}^{d,k}(a)$.

Example 1. Let us look at the special case $f = x^m + 1$ and m = d = k = n. In this case, the $\mathsf{Rot}_f^d(a)$ and $\mathsf{Toep}^{d,k}(a) \cdot \mathsf{Rot}_f^{d+k-1}(1)$ matrices are the following ones:

$$\operatorname{Foep}^{n,n}(a) \cdot \operatorname{Rot}_{f}^{2n-1}(1) = \begin{pmatrix} a_{0} & a_{1} & \dots & a_{n-2} & a_{n-1} \\ -a_{n-1} & a_{0} & \dots & a_{n-3} & a_{n-2} \\ \vdots & \ddots & & \vdots \\ -a_{1} & -a_{2} & \dots & -a_{n-1} & a_{0} \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ 0 & a_{0} & \dots & a_{n-2} & a_{n-1} & \dots & 0 \\ \vdots & \ddots & & \ddots & & \\ \vdots & \ddots & & \ddots & & \\ 0 & 0 & \dots & \dots & a_{0} & \dots & a_{n-2} & a_{n-1} \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \ddots & & \vdots \\ 0 & 0 & \dots & 0 & 1 \\ -1 & 0 & \dots & 0 & 0 \\ 0 & -1 & \dots & 0 & 0 \\ \vdots & \ddots & & \vdots \\ 0 & 0 & \dots & -1 & 0 \end{pmatrix}$$

Moreover, in this case, for any polynomial $b \in R^{<n}[x]$, the vector $\operatorname{Rot}_f(a) \cdot \overline{\mathbf{b}}$ is exactly $\overline{ab \mod f}$, which implies that the vector of coefficients of $\overline{ab \mod f}$ can be written as a product of $\operatorname{Toep}^{n,n}(a)$ and $\operatorname{Rot}_f^{2n-1}(1) \cdot \overline{\mathbf{b}}$.

The expansion factor of a polynomial f was introduced in [LM06] and measures how large become the coefficients of a polynomial g when it is reduced modulo f, compared to the coefficients of g.

Definition 2.13. Let $f \in \mathbb{Z}[x]$ of degree m. Then the expansion factor of f is defined as $\mathsf{EF}(f) = \max(\|g \mod f\|_{\infty}/\|g\|_{\infty} : g \in \mathbb{Z}^{<2m-1}[x] \setminus \{0\}).$

There is no known polynomial time algorithm to compute the expansion factor of an arbitrary polynomial. Still, some bounds have been obtained in [LM06]. Using these bounds, one can obtain special classes of polynomials whose expansion factors are polynomially bounded.

Definition 2.14. If a polynomial $f(x) \in \mathbb{Z}[x]$ can be written as $f(x) = x^n + \sum_{i=0}^{n-m} f_i x^i$ with $f_{n-m} \neq 0$ and $0 < m \leq n$, we say that gap(f) = m.

Lemma 2.3 (Adapted from [LM06, Th. 3.6]). If f is a polynomial in $\mathbb{Z}[x]$,

$$\mathsf{EF}(f) \le 2||f||_{\infty} \cdot (2||f||_1)^{\lceil \frac{\deg(f)-2}{\operatorname{gap}(f)}\rceil}.$$

As a consequence, for any polynomial $f = x^m + h$, with $h = \sum_{i \le m/2} h_i x^i$ and $||h||_{\infty} \in \mathsf{poly}(m)$, we have that $\mathsf{EF}(f) \in \mathsf{poly}(m)$.

The expansion factor is useful when trying to bound the largest singular value of the matrices \mathbf{M}_{f}^{d} and $\mathsf{Rot}_{f}^{k}(1)$.

Lemma 2.4 ([LVV19, Le. 9]). Let $f(x) \in \mathbb{Z}[x]$ and $d \leq \deg(f)$. Then $\|\mathbf{M}_f^d\| \leq \sqrt{d} \cdot \mathsf{EF}(f)$.

Lemma 2.5. Let $f \in \mathbb{Z}[x]$ and $k \geq \deg(f)$. Then $\|\mathsf{Rot}_f^k(1)\|^2 \leq \deg(f) + (k - \deg(f)) \cdot \deg(f) \cdot \mathsf{EF}(f)^2$.

Proof. The bound on $\|\operatorname{Rot}_{f}^{k}(1)\|$ can be obtained by noticing that $\operatorname{Rot}_{f}^{k}(1)$ contains $\mathbf{I}_{\deg(f)}$ as a submatrix and all its other entries are bounded by $\mathsf{EF}(f)$.

2.3 Probabilities

We recall now the probablity concepts that we are going to use throughout this thesis.

2.3.1 Basic definitions

In this section, we recall the definitions of the statistical distance, the Rényi divergence and the minentropy. The statistical distance and the Rényi divergence of two distributions are both measures of similarity of the respective distributions, while the min-entropy of a random variable measures the probability of the most likely result.

Definition 2.15. If D_1, D_2 are two continuous distributions over the same measurable set S, their statistical distance is $\Delta(D_1, D_2) := \int_S |D_1(x) - D_2(x)| \, dx$. If D_1 and D_2 are defined over some finite set S, their statistical distance is $\Delta(D_1, D_2) := \frac{1}{2} \sum_{x \in S} |\Pr[D_1 = x] - \Pr[D_2 = x]|$.

Definition 2.16. The Rényi divergence of two continuous distributions D_1 and D_2 over the same measurable set S is $R(D_1||D_2) = \int_S D_1(x)^2/D_2(x) dx$.

Definition 2.17. Let X be a random variable chosen according to a discrete distribution D defined on a set S. We define the min-entropy of X as $H_{\infty}(X) = -\max_{x \in S} \log \Pr[X = x]$.

2.3.2 Leftover hash lemma

We are going to use the following variant of the leftover hash lemma borrowed from [HILL99] in Chapter 4 in order to approximate how far a certain distribution is from the uniform one. In Chapter 3, we will prove a variant of the leftover hash lemma over number rings for specific distributions.

Definition 2.18. A finite family \mathcal{H} of hash functions $h: X \to Y$ is called universal if $\Pr_{h \leftrightarrow U(\mathcal{H})}[h(x_1) = h(x_2)] = 1/|Y|$, for all $x_1 \neq x_2 \in X$.

Lemma 2.6 (Leftover hash lemma). Let X, Y, Z denote finite sets. Let \mathcal{H} be a universal family of hash functions $h: X \to Y$. Let $f: X \to Z$ be arbitrary. Then for any random variable T taking values in X, we have:

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

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

2.3.3 Gaussian distributions

In the problems that we will study, the so-called noise distributions will be Gaussian. In this section, we recall the definition of a Gaussian distribution.

A symmetric matrix $\Sigma \in \mathbb{R}^{n \times n}$ is positive definite if $x^t \Sigma x > 0$ for every non-zero vector $x \in \mathbb{R}^n$. For any non-singular matrix $\mathsf{B} \in \mathbb{R}^{n \times n}$, the matrix $\Sigma = \mathsf{B}\mathsf{B}^t$ is positive definite and we say that $\mathsf{B} = \sqrt{\Sigma}$. Every positive definite matrix Σ has a square root $\mathsf{B} = \mathsf{Q}\mathsf{D}$, where $\Sigma = \mathsf{Q}\mathsf{D}^2\mathsf{Q}^t$ is the spectral decomposition of Σ . Note that the square root of a positive definite matrix is not unique ($\mathsf{B}' = \mathsf{B}\mathsf{H}$ is also a square root of Σ for every orthogonal matrix $\mathsf{H} \in \mathbb{R}^{n \times n}$). If $\Sigma \in \mathbb{R}^{n \times n}$ is a positive definite matrix, its inverse is also positive definite and, moreover, the set of positive definite matrices is closed under addition.

Definition 2.19. For a positive definite matrix $\Sigma \in \mathbb{R}^{n \times n}$, we define the Gaussian function on \mathbb{R}^n of covariance matrix Σ as $\rho_{\Sigma}(x) = \exp(-\pi x^t \Sigma^{-1} x)$ for every $x \in \mathbb{R}^n$. The probability distribution whose density is proportional to ρ_{Σ} is called the Gaussian distribution and is denoted D_{Σ} .

When $\Sigma = \operatorname{diag}(s_i^2)_i$ for some $\mathbf{s} = (s_1, \ldots, s_n)^t \in \mathbb{R}^n$, we write $\rho_{\mathbf{s}}$ and $D_{\mathbf{s}}$ instead of ρ_{Σ} and D_{Σ} , respectively. When $\Sigma = s^2 \cdot \mathbf{I}_n$ for some $s \in \mathbb{R}$, we simply write ρ_s and D_s .

2.3.4 Gaussian distributions over lattices

We now define Gaussian distributions on lattices, which we call from now on *discrete* as lattices represent discrete subgroups of \mathbb{R}^n . We then show that if the standard deviation of a discrete Gaussian r is large enough, the distribution behaves similarly to a continuous one.

Definition 2.20. For a positive definite matrix $\Sigma \in \mathbb{R}^{n \times n}$ and a full-rank lattice $L \subset \mathbb{R}^n$ we define $\rho_{\Sigma}(L) := \sum_{x \in L} \rho_{\Sigma}(x)$. Using this, we can now define the discrete Gaussian distribution over L of covariance parameter Σ as $D_{L,\Sigma}(x) = \rho_{\Sigma}(x)/\rho_{\Sigma}(L)$ for every $x \in L$.

In Figure 2.4, we plot a Gaussian distribution on \mathbb{Z} of standard deviation 1.2.

Lemma 2.7 ([LPSS14, Le. 5]). Let $\Sigma_1, \Sigma_2 \in \mathbb{R}^{n \times n}$ two covariance matrices and $\mathsf{L}_1, \mathsf{L}_2$ full-rank lattices in \mathbb{R}^n such that $1 \ge \eta_{\epsilon}((\Sigma_1^{-1} + \Sigma_2^{-1})^{1/2} \cdot (\mathsf{L}_1 \cap \mathsf{L}_2))$ for some $\epsilon \in (0, 1/2)$. If $x_1 \leftrightarrow D_{\mathsf{L}_1, \Sigma_1}$ and $x_2 \leftrightarrow D_{\mathsf{L}_2, \Sigma_2}$, then the statistical distance between the distribution of $x_1 + x_2$ and $D_{\mathsf{L}_1 + \mathsf{L}_2, \Sigma_1 + \Sigma_2}$ is less than 4ϵ .

Lemma 2.8 ([Ban95, Le. 2.10]). Let L be a full-rank lattice in \mathbb{R}^n and r > 0. Then $\Pr_{\mathbf{x} \leftarrow D_{\mathsf{L},r}}(\|\mathbf{x}\|_{\infty} > r \cdot t) \leq 2n \cdot \exp(-\pi \cdot t^2)$.

The smoothing parameter of a lattice L was introduced in [MR04] and it informally says how large the parameter r should be in order for the distribution $D_{L,r}$ to behave similarly to the continuous Gaussian D_r .

Definition 2.21. For $\epsilon > 0$, we define the smoothing parameter $\eta_{\epsilon}(\mathsf{L})$ as the smallest r > 0 such that $\rho_{1/r}(\mathsf{L}^* \setminus \{0\}) \leq \epsilon$.

If $L_1 \subseteq L_2$ are two lattices, we have that $\eta_{\epsilon}(L_2) \leq \eta_{\epsilon}(L_1)$ for any $\epsilon > 0$.

Lemma 2.9 ([MR04, Le. 3.3]). For any full-rank lattice $L \subset \mathbb{R}^n$ and $\epsilon > 0$, we have $\eta_{\epsilon}(L) \leq \lambda_n(L) \cdot \sqrt{\ln(2n(1+1/\epsilon))/\pi}$.

Lemma 2.10 (Adapted from [MR04, Le. 4.4]). Let L be an *n*-dimensional lattice, $\epsilon \in (0, 1/3)$ and $r \geq \eta_{\epsilon}(L)$. Then $\Pr_{\mathbf{x} \leftarrow D_{L,r}}[\|\mathbf{x}\| \geq 2r\sqrt{n}] \leq 2^{-2n}$.

Lemma 2.11 ([GPV08, Cor. 2.8]). Let $\mathsf{L}' \subseteq \mathsf{L}$ be full-rank lattices, $\epsilon \in (0, 1/2)$ and $r \geq \eta_{\epsilon}(\mathsf{L}')$. Then $\Delta(D_{\mathsf{L},r} \mod \mathsf{L}', U(\mathsf{L}/\mathsf{L}')) \leq 2\epsilon$.

Lemma 2.12 ([PR06, Le. 2.11]). Let L be an n-dimensional lattice, $\epsilon \in (0, 1/3)$ and $r \ge 4\eta_{\epsilon}(L)$. Then $D_{L,r}(0) \le 2^{-2n+1}$.



Fig. 2.4: Discrete Gaussian distribution on \mathbb{Z} of standard deviation r = 1.2.

2.4 Algebraic number theory

In this section, we give a brief introduction into algebraic number theory. We reccommend [Ste17] for a more general description of the results.

2.4.1 Number fields

A number field K is a finite extension of \mathbb{Q} , which can always be described as $\mathbb{Q}[x]/f$ for some monic irreducible polynomial $f \in \mathbb{Z}[x]$, or $\mathbb{Q}[\alpha]$ for some root α of f. Note that a given K admits several such f's. In this setup, the polynomial f is called a *defining polynomial* of K and the *extension degree* of K is deg f. The set of all elements of K whose minimal polynomials have coefficients in \mathbb{Z} is a ring called the *ring of integers* and is denoted by \mathcal{O}_K . It contains the subring $\mathbb{Z}[\alpha] \simeq \mathbb{Z}[x]/f$ and, in general, the inclusion is strict. Examples where $\mathcal{O}_K = \mathbb{Z}[\alpha]$ include quadratic extensions $\mathbb{Q}(\sqrt{d})$ for $d \equiv 2, 3$ mod 4, cyclotomic fields (*i.e.*, when α is a primitive root of the unity) and number fields with a defining polynomial f of squarefree discriminant Δ_f . To avoid confusion with elements of \mathcal{O}_K , elements in \mathbb{Z} are called *rational integers*.

2.4.2 Embeddings

A number field $K = \mathbb{Q}[\alpha]$ of degree *n* has exactly *n* ring embeddings $\sigma_i : K \to \mathbb{C}$ in the complex field. If we let $\alpha_1, \ldots, \alpha_n$ be the *n* roots of its defining polynomial, then these embeddings are defined by $\sigma_i(\alpha) = \alpha_i$ and extended \mathbb{Q} -linearly. They are often called *Minkowski embeddings*. If the image of an embedding is contained in the real field \mathbb{R} it is said to be *real*, else it is said to be *complex*. As complex roots come by pairs of conjugates, so do the complex embeddings. We let s_1 denote the number of real embeddings and s_2 the number of pairs of complex embeddings, so that $n = s_1 + 2s_2$.

We define the canonical space $H := \{ \mathbf{x} \in \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2} : \forall i \leq s_2 : x_{s_1+s_2+i} = \overline{x_{s_1+i}} \}$. The space H with the inner product induced on it from \mathbb{C}^n is isomorphic to \mathbb{R}^n as inner product spaces. Indeed, let $(\mathbf{e}_i)_{i\leq n}$ be the canonical basis of \mathbb{C}^n . We define $\mathbf{h}_i = \mathbf{e}_i$ for $i \leq s_1$, and $\mathbf{h}_{s_1+i} = (\mathbf{e}_{s_1+i} + \mathbf{e}_{s_1+s_2+i})/\sqrt{2}$ and $\mathbf{h}_{s_1+s_2+i} = (\mathbf{e}_{s_1+i} - \mathbf{e}_{s_1+s_2+i})/\sqrt{-2}$ for $i \leq s_2$. The \mathbf{h}_i 's form an orthonormal \mathbb{R} -basis of H. The embedding map, which is usually called *canonical* or *Minkowski*, is then defined as $\sigma : K \to H$ by mapping an element in K to its vector of (suitably ordered) embeddings. Note that via the embedding map, we have $K_{\mathbb{R}} := K \otimes_{\mathbb{Q}} \mathbb{R} \simeq H$. Among its nice properties, the multiplicative structure of K is preserved, *i.e.*, $\sigma(xy) = (\sigma_1(x)\sigma_1(y), \ldots, \sigma_n(x)\sigma_n(y))$.

If we are given a (geometric) norm $\|\cdot\|$ on the space $\mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$, then we can consider the geometric norm of an element in K by means of the Minkowski embeddings. The (field) *trace* is the \mathbb{Q} -linear map defined as $\operatorname{Tr}(x) = \sum_{i < n} \sigma_i(x)$ and the (field) *norm* is $N(x) = \prod_{i < n} \sigma_i(x)$.

Another way is to use the so-called *coefficients* embedding, which amounts to viewing an element $a(x) = \sum_{i=0}^{n} a_i x^i$ as its vector of coefficients $\mathbf{a} = (a_i)_{i < n}$. Different defining polynomials for $K = \mathbb{Q}[x]/f$ give different coefficient embeddings, and coefficient and Minkowski embeddings have different geometric settings. Going from the coefficient representation \mathbf{a} of K to its Minkowski equivalent is done by the linear transformation $\sigma(a) = V_f \mathbf{a}$, where V_f denotes the Vandermonde matrix of $f = \prod_{i=1}^{n} (x - \alpha_i)$:

$$V_f = \begin{pmatrix} 1 & \alpha_1 & \dots & \alpha_1^{n-1} \\ 1 & \alpha_2 & \dots & \alpha_2^{n-1} \\ \vdots & & \dots & \vdots \\ 1 & \alpha_n & \dots & \alpha_n^{n-1} \end{pmatrix}.$$

It is well-known that the square determinant of this matrix is the discriminant of f, *i.e.*, we have $(\det V_f)^2 = \Delta_f = \prod_{i \neq j} (\alpha_i - \alpha_j)$. When it defines a number field, the polynomial f does not have any double root thus V_f is invertible and we have $\mathbf{a} = V_f^{-1}\sigma(a)$.

2.4.3 Rings and ideals in number fields

We call any subring of K a number ring. For a number ring R, an (integral) R-ideal is an additive subgroup $I \subseteq R$ which is closed by multiplication in R, *i.e.*, such that IR = I. A more compact definition is to say that I is an R-module contained in R. If a_1, \ldots, a_k are elements in R, we let $\langle a_1, \ldots, a_k \rangle = a_1R + \ldots + a_kR$ and call it the ideal generated by the a_i 's. The product of two ideals I, Jis the ideal generated by all elements xy with $x \in I$ and $y \in J$. The sum, product and intersection of two R-ideals are again R-ideals.

Two integral *R*-ideals *I*, *J* are said to be coprime if I + J = R, and, in this case, we have $I \cap J = IJ$. Any non-zero ideal in a number ring has finite index, *i.e.*, the quotient ring R/I is always finite when *I* is a non-zero *R*-ideal. An *R*-ideal \mathfrak{p} is said to be prime if whenever $\mathfrak{p} = IJ$ for some *R*-ideals *I*, *J*, then either $I = \mathfrak{p}$ or $J = \mathfrak{p}$. In a number ring, any prime ideal \mathfrak{p} is maximal, *i.e.*, *R* is the only *R*-ideal containing it. It also means that the quotient ring R/\mathfrak{p} is a finite field. It is well-known that any \mathcal{O}_K -ideal admits a unique factorization into prime \mathcal{O}_K -ideals, *i.e.*, it can be written $I = \mathfrak{p}_1^{e_1} \dots \mathfrak{p}_k^{e_k}$ with all \mathfrak{p}_i 's distinct prime ideals. It fails to hold in general number rings and orders, but we describe later in Lemma 3.1 how the result can be extended in certain cases.

A fractional *R*-ideal *I* is an *R*-module such that $xI \subseteq R$ for some $x \in K^{\times}$. An integral ideal is a fractional ideal, and so are the sum, the product and the intersection of two fractional ideals. A fractional *R*-ideal *I* is said to be invertible if there exists a fractional *R*-ideal *J* such that IJ = R. In this case, the (unique) inverse is the integral ideal $I^{-1} = \{x \in K : xI \subseteq R\}$. Any \mathcal{O}_K -ideal is invertible, but it is again false for a general number ring.

The algebraic norm of a non-zero integral *R*-ideal *I* is defined as $\mathcal{N}_R(I) = |R/I|$, and we will omit the subscript when $R = \mathcal{O}_K$. It satisfies $\mathcal{N}_R(IJ) = \mathcal{N}_R(I)\mathcal{N}_R(J)$ for every *R*-ideals *I*, *J*.

The dual of a fractional *R*-ideal *I* is $I^{\vee} = \{\alpha \in K : \operatorname{Tr}(\alpha I) \subseteq \mathbb{Z}\}$, which is also a fractional *R*-ideal. We always have $II^{\vee} = R^{\vee}$, so that $I^{\vee} = I^{-1}R^{\vee}$ when *I* is invertible. We also have $I^{\vee\vee} = I$ for any *R*-ideal *I*.

A particularly interesting dual is \mathcal{O}_K^{\vee} , whose inverse $(\mathcal{O}_K^{\vee})^{-1}$ is called the different ideal. The different ideal is an integral ideal, whose norm $\Delta_K = \mathcal{N}((\mathcal{O}_K^{\vee})^{-1})$ is called the discriminant of the number field. We note that, for every f defining K, the field discriminant Δ_K is a factor of the discriminant Δ_f of f. This provides an upper bound on Δ_K in terms of the defining polynomial f.

2.4.4 Ideal lattices

A lattice could also be defined as a full-rank discrete additive subgroup of an \mathbb{R} -vector space V which is a Cartesian power H^m (for $m \geq 1$) of H. Any fractional \mathcal{O}_K -ideal I is a free \mathbb{Z} -module of rank $n = \deg(K)$, *i.e.*, it can be written as $\mathbb{Z}u_1 + \cdots + \mathbb{Z}u_n$ for some u_i 's in K. Its canonical embedding $\sigma(I)$ is a lattice of dimension n in the \mathbb{R} -vector space $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$. Such a lattice is called an ideal lattice (for \mathcal{O}_K). For the sake of readability, we will abuse notations and often identify I and $\sigma(I)$. It is possible to look at the coefficient embedding of such lattices as well, but we will not need it. The lattice corresponding to I^{\vee} is $\overline{I^*}$. The discriminant of K satisfies $\Delta_K = (\det \mathcal{O}_K)^2$. In the following lemma, the upper bounds follow from Minkowski's theorem whereas the lower bounds are a consequence of the algebraic structure underlying ideal lattices.

Lemma 2.13 (Adapted from [PR07, Se. 6.1]). Let K be a number field of degree n. For any fractional \mathcal{O}_K -ideal I, we have:

$$\begin{array}{rcl} \sqrt{n} \cdot \mathcal{N}(I)^{1/n} & \leq & \lambda_1(I) & \leq & \sqrt{n} \cdot (\mathcal{N}(I)\sqrt{\Delta_K})^{1/n}, \\ \mathcal{N}(I)^{1/n} & \leq & \lambda_1^{\infty}(I) & \leq & (\mathcal{N}(I)\sqrt{\Delta_K})^{1/n}. \end{array}$$

Lemma 2.14 (Adapted from [PR07, Le. 6.5]). For any \mathcal{O}_K -ideal I and $\epsilon \in (0,1)$, we have $\eta_{\epsilon}(I) \leq \sqrt{\log(2n(1+1/\epsilon))/(\pi n)} \cdot (\mathcal{N}(I)\Delta_K)^{1/n}$.

2.4.5 Gaussians on H

We define the Gaussian distribution D_{Σ}^{H} on H as the distribution obtained by sampling $x \leftrightarrow D_{\Sigma}$ and returning $\sum_{i} x_{i} \mathbf{h}_{i}$. The canonical embedding allows us to interpret a distribution on H as a distribution on K. We will repeatedly use the observation that if \mathbf{x} is sampled from D_{Σ}^{H} and t belongs to $K_{\mathbb{R}}$, then $t \cdot \mathbf{x}$ is distributed as $D_{\Sigma'}^{H}$ with $\Sigma' = \operatorname{diag}(|\sigma_{i}(t)|) \cdot \Sigma \cdot \operatorname{diag}(|\sigma_{i}(t)|)$.

2.5 Cryptographic definitions

In this section, we introduce the most relevant cryptographic concepts that we are going to use throughout the thesis.

2.5.1 Proofs by reduction

In modern cryptography, if we wish to prove that a cryptographic scheme is secure, we rely on some presumably hard problem (GapSVP, GapCVP, etc.). Concretely, we prove the security by *reduction*, *i.e.*, we show how to transform any efficient adversary \mathcal{A} that succeeds in breaking the cryptographic scheme Π in time t with advantage ϵ into an efficient algorithm \mathcal{B} that solves the hard problem P in time t' with probability ϵ' . If $t \approx t'$ and $\epsilon \approx \epsilon'$, we say that the reduction is *tight*. When $t' \gg t$ or $\epsilon' \ll \epsilon$ we call the reduction *non-tight*. If the parameters of Π are set based on the concrete hardness assumption for the problem P and the reduction from P to Π is non-tight, the parameters will be larger than in the case of a tight reduction.

2.5.2 (Quantum) Random-oracle model

The random-oracle model is a general framework which allows the construction of cryptographic schemes and the proof of their security assuming the existence of efficient truly random functions. More specifically, the random-oracle model assumes the existence of a public random function H which can be evaluated by querying an oracle. The oracle returns a uniform value H(x) when given an input x which has not been previously queried and is consistent with its previous answer when given an input x which has already been queried. When the cryptographic primitive is implemented in practice, the truly random oracle H is replaced by a cryptographic hash function. Although there is no theoretical evidence that a security proof in the random-oracle model still holds when the oracle is instantiated with a specific cryptographic hash function, the random-oracle model could be seen as a tradeoff between a rigorous proof of security and no proof.

In the quantum random-oracle model, the oracle H can be queried on quantum superpositions. The cryptographic definitions and the security games in the quantum random-oracle model are similar to those in the classical random-oracle model, with the only difference that they make use of an adversary which is given quantum access to the random oracles involved, but classical access to all the other oracles (e.g. signing/decryption oracles). We call any such adversary quantum. The framework where we do not assume the existence of truly random functions is called the *standard model*. For a gentle introduction to the random-oracle model we recommend [KL14], while for a background on quantum computing, we recommend [dW19].

2.5.3 Pseudorandom functions

In this section we give the formal definition of pseudorandom functions.

Definition 2.22. A pseudorandom function PRF is a map $PRF : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^k$ where \mathcal{K} is a finite key space and n, k are positive integers, such that no (quantum) adversary A trying to distinguish the output of the PRF from a uniform output is able to have a non-negligible advantage. The adversary A is able to perform a polynomial number of classical queries to one of the two oracles $PRF(K, \cdot)$ and $RF(\cdot)$ and its advantage is defined as

$$\operatorname{Adv}_{\mathsf{PRF}}^{\mathsf{PR}}(A) := |\operatorname{Pr}(A^{\mathsf{PRF}(K,\cdot)} = 1 | K \leftarrow \mathcal{K}) - \operatorname{Pr}(A^{\mathsf{RF}(\cdot)} = 1)|$$

where $\mathsf{RF}: \{0,1\}^n \to \{0,1\}^k$ is a uniformly sampled function from $\mathcal{F}(\{0,1\}^n, \{0,1\}^k)$.

The evaluation in \mathbf{x} of $\mathsf{PRF}(K, \cdot)$ is deterministically computable in polynomial time for any $K \in \mathcal{K}$ and any valid input \mathbf{x} . Defining the function RF takes exponential time, since the domain of the function is exponentially large. Still, since the adversary has access only to a polynomial number of evaluations of RF, we can consider that its defining the function takes only polynomial time: when the adversary asks for the evaluation of RF on an input x, we either pick the output $\mathsf{RF}(x)$ uniformly at random in $\{0,1\}^k$ and store the pair $(x, \mathsf{RF}(x))$ in a table if x has never been queried before, or, if x has already been queried, we return the value y such that (x, y) is in this table.

2.5.4 Public-key encryption

Public-key encryption allows two parties to confidentially communicate without sharing a common key before they interact. In this section, we formalize this notion and recall the related security definitions that we are going to use in Chapter 5.

Definition 2.23 (Public-key encryption scheme). A public-key encryption scheme is a tuple of classical ppt (i.e. probabilistic polynomial time) algorithms PKE := (KeyGen, Enc, Dec).

- The key generation algorithm KeyGen takes as input a security parameter λ in unary and returns the public key pk and the secret key sk. The public key pk determines the set of messages M and the set of ciphertexts C.
- The encryption algorithm Enc takes as input the public key pk and a message m ∈ M and returns the ciphertext c := Enc_{pk}(m) ∈ C.

- The challenger runs KeyGen, gets the pair (pk,sk) and publishes pk.
- After having been given the public key pk, the adversary A outputs two messages m_0 and m_1 .
- The challenger chooses a uniform bit $b \in \{0, 1\}$ and sends back to the adversary $c := \mathsf{Enc}_{\mathsf{pk}}(m_b)$.
- A outputs $b' \in \{0, 1\}$.

The output of this experiment is 1 if b = b' and 0 if $b \neq b'$.

Fig. 2.5: The $\mathsf{Pub}_{\mathsf{PKE}}^{\mathsf{CPA},A}$ experiment.

- The challenger runs KeyGen, gets the pair (pk, sk) and publishes pk.
- After having been given the public key pk, the adversary may query $\mathsf{Dec}_{\mathsf{sk}}(\cdot)$ many times.
- The adversary A outputs two messages m_0 and m_1 .
- The challenger chooses a uniform bit $b \in \{0, 1\}$ and sends back to the adversary $c := \mathsf{Enc}_{\mathsf{pk}}(m_b)$.
- The adversary can still interact with the decryption oracle, but may not ask for a decryption of the ciphertext c.
- A outputs $b' \in \{0, 1\}$.

The output of this experiment is 1 if b = b' and 0 if $b \neq b'$.

Fig. 2.6: The $\mathsf{Pub}_{\mathsf{PKE}}^{\mathsf{CCA},A}$ experiment.

 The decryption algorithm Dec takes as input the secret key sk and a ciphertext c and returns Dec_{sk}(c) ∈ M ∪ {⊥}.

We say that \mathcal{M} is the message space, \mathcal{C} is the ciphertext space and \perp is the failure symbol.

Definition 2.24. We say that the public-key encryption scheme PKE has correctness error $\delta \ge 0$ if for any message $m \in \mathcal{M}$,

$$\Pr[\mathsf{Dec}_{\mathsf{sk}}(\mathsf{Enc}_{\mathsf{pk}}(m)) = m] > 1 - \delta,$$

where the probability is taken over the output (pk, sk) of KeyGen and the randomness used by the Enc and Dec algorithms.

The security of PKE is modelled by different games. We recall here two of the most relevant security notions related to public-key encryption. Intuitively, we say that PKE is indistinguishable against chosen-plaintext attacks (IND-CPA secure) if no efficient adversary can recognize which of two messages is encrypted in a given ciphertext, even if the two candidate messages have been chosen by itself. Formally, given a public-key encryption scheme PKE := (KeyGen, Enc, Dec) and an adversary A, first consider the four steps experiment $Pub_{PKE}^{CPA,A}$ captured in Figure 2.5.

Definition 2.25. We say that PKE is indistinguishable against chosen-plaintext attacks (IND-CPA secure) if for all ppt (quantum) adversaries A, there is a negligible function negl such that

$$\Pr[\mathsf{Pub}_{\mathsf{PKE}}^{\mathsf{CPA},A} = 1] < \frac{1}{2} + \mathsf{negl}(\lambda).$$

By allowing the adversary A to query a decryption oracle $\mathsf{Dec}_{\mathsf{sk}}(\cdot)$, we can define a different experiment $\mathsf{Pub}_{\mathsf{PKE}}^{\mathsf{CCA},A}$ as in Figure 2.6.

Definition 2.26. We say that PKE is indistinguishable against chosen-ciphertext attacks (IND-CCA secure) if for all ppt (quantum) adversaries A, there is a negligible function negl such that

$$\Pr[\mathsf{Pub}_{\mathsf{PKE}}^{\mathrm{CCA},A} = 1] < \frac{1}{2} + \mathsf{negl}(\lambda).$$

It is usually harder to prove directly the IND-CCA security of a public-key encryption scheme. Still, there exist several transformations that turn a public-key encryption scheme with weaker security properties into an IND-CCA secure one, both in the random-oracle model ([FO99, FO13, OP01, CHJ⁺02], among others) and quantum random-oracle model ([EU16, HHK17], etc.). The most well-known such transformation, due to Fusjisaki and Okamoto, combines an IND-CPA public-key encryption scheme with a one-time secure symmetric encryption scheme and two hash functions into an encryption scheme that is IND-CCA secure in the quantum random-oracle model.

2.5.5 Digital signatures

Digital signatures allow a signer who has published a verification key vk to sign a message using the corresponding private key sk in such a way that anyone who knows vk can verify that the message originated from the respective signer and whether the message was altered or not. In this section we recall the formal definition of a digital signature, the security notion we are interested in and the Fiat-Shamir transform [FS86]. The Fiat-Shamir transform combines an identification scheme ID := (IGen, P, V) and a hash function into a digital signature scheme SIG := (G = IGen, S, \overline{V}) which is secure in the random-oracle model. We recall that there also exist lattice-based digital signatures in the literature whose security proof holds in the standard model ([CHKP10, DM14], among others).

Definition 2.27 (Digital signature). A digital signature scheme is a tuple of classical ppt algorithms SIG := (KeyGen, Sign, Ver).

- The key generation algorithm KeyGen takes as input a security parameter λ (in unary) and returns the verification key vk and the signing key sk. The two keys determine the set of messages M and the set of possible signatures Σ.
- The signing algorithm Sign takes as input the key sk and a message $m \in \mathcal{M}$ and returns a signature $\sigma = \text{Sign}(\text{sk}, m) \in \Sigma$.
- The verification algorithm Ver takes as input the verification key vk, a message m and a signature σ and returns Ver(vk, m, σ) ∈ {0, 1}.

We say that \mathcal{M} is the message space and Σ is the signature space.

Definition 2.28. We say that the signature scheme SIG has correctness error $\delta \ge 0$ if for any message $m \in \mathcal{M}$,

$$\Pr[\mathsf{Ver}(\mathsf{vk}, m, \mathsf{Sign}(\mathsf{sk}, m)) = 1] > 1 - \delta_2$$

where the probability is taken over the output (sk, vk) of KeyGen and the randomness used by the Ver and Sign algorithms.

Definition 2.29 (Unforgeability). A signature scheme SIG := (KeyGen, Sign, Ver) is said to be unforgeable against one-per-message chosen message attack (UF-CMA₁) in the (quantum) random oracle model if for every ppt (quantum) forger \mathcal{F} , there is a negligible function negl such that

 $\Pr[$ the output of the experiment in Figure 2.7 is 1 $] < \operatorname{negl}(\lambda),$
- The challenger runs KeyGen, gets the pair (vk, sk) and publishes vk.
- The forger \mathcal{F} may adaptively (quantumly) query the random oracle $H(\cdot)$ and classically query the signing oracle $\mathsf{Sign}_{\mathsf{sk}}(\cdot)$ a polynomial number of times in order to obtain signatures for the distinct messages $\{m_i\}_i$.
- The forger \mathcal{F} outputs a message $m^* \neq m_i$ for all i and σ^* .

The output of the experiment is 1 if $Ver(vk, m^*, \sigma^*) = 1$ and 0 otherwise.

Fig. 2.7: The $\mathsf{UF}\text{-}\mathsf{CMA}_1$ experiment.

where the probability is taken over the randomness used by KeyGen, Sign, Ver and \mathcal{F} and is denoted by $\operatorname{Adv}_{SIG}^{UF-CMA_1}(\mathcal{F})$.

In the UF-CMA₁ experiment, the forger is not allowed to ask for more than one signature per message. One can extend this definition to the scenario where the attacker may have access to more than one signature for any of poly(n) adaptively chosen messages $\{m_i\}$. In that case, if no (quantum) adversary \mathcal{F} can produce a valid signature for a message $m^* \notin \{m_i\}$, we say that the signature scheme is unforgeable against chosen message attack (UF-CMA) in the (quantum) random-oracle model.

As showed in [BBS16], a UF-CMA₁ signature scheme can be combined with a pseudo-random function to obtain a signature scheme that is UF-CMA, and the conversion is tight in the random-oracle model (further, the upgrade preserves strongness). As observed in [KLS18], this transformation is also tight in the quantum random-oracle model.

In the corresponding *strong* UF-CMA/UF-CMA₁ experiments, the adversary may return a forgery for a message which has already been queried to the signing oracle, but with a different signature. As shown in [Kat10, p. 27], any UF-CMA secure signature can be upgraded to a UF-sCMA secure signature using a one-time UF-sCMA secure signature ([Lam79, BDE⁺11], etc.).

2.5.5.1 Identification schemes

A canonical identification scheme is an interactive protocol between two parties: a prover P and a verifier V that allows P to prove its identity (*i.e.*, to authenticate itself) to V. The prover sends a commitment W and the verifier selects a uniform challenge c and sends it to P. Upon receiving c, the prover sends back a response Z to the verifier. After it receives Z, the verifier makes a deterministic decision. In this section we give recall the formal definition and the basic security properties of an identification scheme. We closely follow the notations used in [KLS18].

Definition 2.30 (Canonical identification scheme). A canonical identification scheme is a tuple of classical ppt algorithms ID := (IGen, P, V).

- The key generation algorithm IGen takes as input a security parameter λ and returns the public and secret keys (pk, sk). The public key defines the set of challenges ChSet, the set of commitments WSet, and the set of responses ZSet.
- The prover algorithm P consists of two sub-algorithms: P₁ takes as input the secret key sk and returns a commitment W ∈ WSet and a state St; P₂ takes as inputs the secret key sk, a commitment W, a challenge c, and a state St and returns a response Z ∈ ZSet ∪ {⊥}, where ⊥∉ ZSet is a special symbol indicating failure.
- The verifier algorithm V takes as inputs the public key pk and the conversation transcript (W, c, Z) and outputs 1 (acceptance) or 0 (rejection).

If $Z = \bot$, then we set $(W, c, Z) = (\bot, \bot, \bot)$.

The triple $(W, c, Z) \in \mathsf{WSet} \times \mathsf{ChSet} \times \mathsf{ZSet} \cup \{(\bot, \bot, \bot)\}$ generated in this way is called a *transcript*. Given the public key pk, the transcript is valid if V(pk, W, c, Z) = 1.

We say that ID has correctness error δ if for all public and secret keys generated by IGen, all possible transcripts in WSet × ChSet × ZSet with $Z \neq \bot$ are valid and the probability that a honestly generated transcript is (\bot, \bot, \bot) is less than δ .

We recall now the background on identification schemes necessary to understand the statement of Theorem 2.3.

Definition 2.31. We say that the canonical identification scheme ID has α bits of min-entropy if

$$\Pr_{(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{IGen}(\lambda)}\left(H_{\infty}(W|(W,St)\leftarrow P_1(\mathsf{sk})\right)\geq \alpha\right)\geq 1-2^{-\alpha}.$$

Definition 2.32 (No-abort honest-verifier zero-knowledge). A canonical identification scheme ID is ε_{zk} -perfect no-abort honest-verifier zero-knowledge (ε_{zk} -perfect na-HVZK) if there exists a ppt algorithm Sim which given only the public key pk outputs (W, c, Z) such that the statistical distance between $(W, c, Z) \leftarrow \text{Sim}(pk)$ and $(W, c, Z) \leftarrow \text{Trans}(pk)$ is at most ε_{zk} and the element c from $(W, c, Z) \leftarrow \text{Sim}(pk)$ follows a uniform distribution conditioned on $c \neq \bot$.

Trans(sk)1: $(W, St) \leftarrow P_1(sk)$ 2: $c \leftarrow ChSet$ 3: $Z \leftarrow P_2(sk, W, c, St)$ 4: if $Z = \perp$ then
5: $(W, c, Z) = (\perp, \perp, \perp)$ 6: end if
7: output (W, c, Z)

Fig. 2.8: The algorithm Trans(sk).

Definition 2.33 (Lossiness). A canonical identification scheme is lossy (and we call it LID) if there exists a lossy key generation algorithm LossylGen that takes as input λ and returns a public key pk_{ls} and no secret key such that the public keys generated by IGen and LossylGen are indistinguishable. In other words, for any quantum adversary A, the following quantity is negligible:

$$Adv_{\mathsf{ID}}^{loss}(A) := |\Pr(A(\mathsf{pk}_{ls}) = 1 | \mathsf{pk}_{ls} \leftarrow \mathsf{LossylGen}(\lambda)) - \Pr(A(\mathsf{pk}) = 1 | (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{IGen}(\lambda))|.$$

Definition 2.34 (Lossy soundness). A canonical identification scheme is ε_{ls} -lossy-sound if, for every quantum adversary A, the following probability that A could impersonate the prover is less than ε_{ls} :

$$\Pr\left[V(\mathsf{pk}_{ls}, W^*, c^*, Z^*) = 1 \middle| \begin{array}{c} \mathsf{pk}_{ls} \leftarrow \mathsf{LossylGen}(\lambda);\\ (W^*, St) \leftarrow A(\mathsf{pk}_{ls});\\ c^* \leftarrow \mathsf{ChSet}; Z^* \leftarrow A(St, c^*) \end{array} \right]$$

2.5.5.2 From identification schemes to digital signatures: the Fiat-Shamir transform

The security of the Fiat-Shamir transform has been analyzed in the quantum random-oracle model in [KLS18, LZ19, DFMS19], among others. The conversion is tight only in [KLS18], but it requires a number of assumptions on the identification scheme. We recall the main result in [KLS18] in the next theorem.

 $\overline{\mathsf{V}}(\mathsf{pk},m,\sigma)$ S(sk,m)1: parse σ as $\sigma = (W, Z)$ 1: i = 02: while $Z = \perp$ and $i \leq k_l$ do 2: c := H(W || m)i = i + 13: output $V(pk, W, c, Z) \in \{0, 1\}$ 3: $(W, St) := \mathsf{P}_1(\mathsf{sk})$ 4: $c := H(W \| m)$ 5: $Z := \mathsf{P}_2(\mathsf{sk}, W, c, St)$ 6: 7: end while 8: if $Z = \perp$ then 9: $\sigma = \perp$ 10: **else** $\sigma = (W, Z)$ 11: 12: end if 13: output σ



Theorem 2.3 ([KLS18, Th. 3.1]). Consider an identification scheme ID which is lossy, ε_{zk} -perfect na-HVZK, has α bits of entropy and is ε_{ls} -lossy sound and the signature scheme SIG obtained by applying the Fiat-Shamir transform to the identification scheme ID, as in Figure 2.9.

For any quantum adversary A against UF-CMA₁ security that issues at most Q_H quantum queries to the random oracle and Q_S classical signing queries, there exists a quantum adversary B against ID such that

$$\operatorname{Adv}_{\mathsf{SIG}}^{\mathsf{UF}-\mathsf{CMA}_1}(A) \leq \operatorname{Adv}_{\mathsf{ID}}^{loss}(B) + 8(Q_H+1)^2 \cdot \varepsilon_{ls} + k_m Q_S \cdot \varepsilon_{zk} + 2^{-\alpha+1}$$

and $Time(B) = Time(A) + k_m Q_H$.

Moreover, if we de-randomize the signature scheme in Figure 2.9 by using a pseudo-random function PRF as in Figure 2.10, then for any quantum adversary A against UF-CMA security that issues at most Q_H quantum queries to the random oracle and Q_S classical signing queries, there exists a quantum adversary B against ID and a quantum adversary C against the PRF such that

$$\operatorname{Adv}_{\mathsf{DSIG}}^{\mathsf{UF-CMA}}(A) \leq \operatorname{Adv}_{\mathsf{ID}}^{loss}(B) + 8(Q_H + 1)^2 \cdot \varepsilon_{ls} + k_m Q_S \cdot \varepsilon_{zk} + 2^{-\alpha+1} + \operatorname{Adv}_{\mathsf{PRF}}^{\mathsf{PR}}(C).$$

The de-randomized version of the signature scheme $\mathsf{DSIG} := (\mathsf{IGen}, \mathsf{DS}, \overline{\mathsf{V}})$ obtained from Fiat-Shamir transformation is given in Figure 2.10. Here, the PRF key K is also a part of the secret key in the signature scheme.



Fig. 2.10: The de-randomized signature DSIG obtained via Fiat-Shamir transform from ID.

CHAPTER 3

ON THE RLWE AND PLWE PROBLEMS

In this chapter, we show that there exist reductions which incur limited parameter losses between the following algebraic variants of LWE: dual-Ring Learning With Errors (dual-RLWE), primal-Ring Learning With Errors (primal-RLWE) and Polynomial Learning With Errors (PLWE), both in their search and decision variants. More precisely: we prove that the (decision/search) dual to primal reduction from Lyubashevsky *et al.* [EUROCRYPT 2010] and Peikert [SCN 2016] can be implemented with a small error rate growth for all rings (the resulting reduction is non-uniform polynomial time); we extend it to polynomial-time reductions between (decision/search) primal RLWE and PLWE that work for a family of polynomials f that is exponentially large as a function of deg f (the resulting reduction is also non-uniform polynomial time); and we exploit the recent technique from Peikert *et al.* [STOC 2017] to obtain a search to decision reduction for RLWE for arbitrary number fields. The reductions incur error rate increases that depend on intrinsic quantities related to K and f.

This chapter is mainly based on a joint work with Damien Stehlé and Alexandre Wallet, published at Eurocrypt 2018.

Contents

3.1	Intr	oduction	40					
3.2 Contributions								
	3.2.1	Techniques	41					
	3.2.2	Related works	42					
	3.2.3	Impact	43					
	3.2.4	Follow-up work	44					
3.3	Ord	ers in number fields	44					
3.4	Form	nal definitions of dual-RLWE, primal-RLWE and PLWE	45					
3.5	From	n dual-RLWE to primal-RLWE	46					
3.6	Con	trolling the noise growth in the dual to primal reduction	47					
3.7	From	n primal-RLWE to PLWE	49					
	3.7.1	Reducing primal-RLWE to $PLWE^{\sigma}$	50					
	3.7.2	Distortion between embeddings	51					
	3.7.3	A family of polynomials with easily computable distortion	51					
	3.7.4	Other "good" families of polynomials	54					
3.8	On s	small elements and $f'(\alpha)$	55					
3.9	Sear	ch to decision dual-RLWE	56					
	3.9.1	A ring-based Leftover Hash Lemma	56					
	3.9.2	Search RLWE to decision RLWE	59					
3.10) On '	Vandermonde matrices and the expansion factor	60					

3.1 Introduction

Ring Learning With Errors (RLWE) was introduced by Lyubashevsky *et al.* in [LPR10], as a means of speeding up cryptographic constructions based on LWE [Reg09]. Let K be a number field, \mathcal{O}_K its ring of integers and $q \geq 2$ a rational integer. The search variant of RLWE with parameters K and qconsists in recovering a secret $s \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ with \mathcal{O}_K^{\vee} denoting the dual of \mathcal{O}_K , from arbitrarily many samples $(a_i, a_i \cdot s + e_i)$. Here each a_i is uniformly sampled in $\mathcal{O}_K/q\mathcal{O}_K$ and each e_i is a small random element of $K_{\mathbb{R}} := K \otimes_{\mathbb{Q}} \mathbb{R}$. The noise term e_i is sampled such that its Minkowski embedding vector follows a Gaussian distribution with a small covariance matrix (relative to $q\mathcal{O}_K^{\vee}$). The decision variant consists in distinguishing arbitrarily many such pairs for a common s chosen uniformly in $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$, from uniform samples in $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$. More formal definitions are provided in Section 3.4, but these suffice for describing our contributions.

Lyubashevsky *et al.* backed in [LPR10] the conjectured hardness of the RLWE problem with a quantum polynomial time reduction from the (worst-case) ApproxSVP restricted to the class of Euclidean lattices corresponding to ideals of \mathcal{O}_K , with geometry inherited from the Minkowski embeddings. They showed its usefulness by describing a public-key encryption with quasi-optimal efficiency: the bit-sizes of the keys and the run-times of all involved algorithms are quasi-linear in the security parameter. A central technical contribution was a reduction from search RLWE to decision RLWE, when K is cyclotomic, and decision RLWE for cyclotomic fields is now pervasive in lattice-based cryptography, including in practice [ADPS16, BDK⁺18, DKL⁺18]. The search-to-decision reduction from [LPR10] was later extended to the case of general Galois rings in [EHL14, CLS19].

Prior to RLWE, Stehlé *et al.* [SSTX09] introduced what is now referred to as Polynomial Ring Learning With Errors (PLWE), for cyclotomic polynomials of degree a power of 2. PLWE is parametrized by a monic irreducible $f \in \mathbb{Z}[x]$ and an integer $q \geq 2$, and consists in recovering a secret $s \in \mathbb{Z}_q[x]/f$ from arbitrarily many samples $(a_i, a_i \cdot s + e_i)$ where each a_i is uniformly sampled in $\mathbb{Z}_q[x]/f$ and each e_i is a small random element of $\mathbb{R}[x]/f$. The decision variant consists in distinguishing arbitrarily many such samples for a common s sampled uniformly in $\mathbb{Z}_q[x]/f$, from uniform samples. Here the noise term e_i is sampled such that its coefficient vector follows a Gaussian distribution with a small covariance matrix. Stehlé *et al.* gave a reduction from the restriction of ApproxSVP to the class of lattices corresponding to ideals of $\mathbb{Z}[x]/f$, to search PLWE, for f a power-of-2 cyclotomic polynomial.

Finally, a variant of RLWE with $s \in \mathcal{O}_K/q\mathcal{O}_K$ rather than $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ was also considered (see, e.g., [DD12] among others), to avoid the complication of having to deal with the dual \mathcal{O}_K^{\vee} of \mathcal{O}_K . In the rest of this thesis, we will refer to the latter as primal-RLWE and to standard RLWE as dual-RLWE.

Even though [LPR10] defined RLWE for arbitrary number fields, the problem was mostly studied in the literature for K cyclotomic. This specialization had three justifications:

- it leads to very efficient cryptographic primitives, in particular if q totally splits over K;
- the hardness result from [LPR10] holds for cyclotomics;
- no particular weakness was known for these fields.

Among cyclotomics, those of order a power of 2 are a popular choice. In the case of a field K defined by the cyclotomic polynomial f, we have that $\mathcal{O}_K = \mathbb{Z}[\alpha]$ for α a root of f. Further, in the case of power-of-2 cyclotomics, mapping the coefficient vector of a polynomial in $\mathbb{Z}[x]/f$ to its Minkowski embedding is a scaled isometry. This makes primal-RLWE and PLWE collapse into a single problem. Still in the case of power-of-2 cyclotomics, the dual \mathcal{O}_K^{\vee} is a scaling of \mathcal{O}_K , implying that dual and primal-RLWE are equivalent. Apart from the monogenicity property, these facts do not hold for all cyclotomics. Nevertheless, Ducas and Durmus [DD12] showed it is still possible to reduce dual-RLWE to primal-RLWE. Later on, Peikert *et. al* [PRSD17] gave a (quantum) reduction from ApproxSVP for \mathcal{O}_K -ideals to decision dual-RLWE for the corresponding field K which works for any number field K.

3.2 Contributions

The focus on the RLWE hardness for non-cyclotomic fields makes the discrepancies between the RLWE and PLWE variants more critical. In this chapter, we show that the six problems considered above — dual-RLWE, primal-RLWE and PLWE, all in both decision and search forms — actually reduce to one another in polynomial time with limited error rate increases, for huge classes of rings. More precisely, these reductions are obtained with the following three results.

- We show that for every field K, it is possible to implement the reduction from decision (resp. search) dual-RLWE to decision (resp. search) primal-RLWE from [LPR10, Le. 2.15] and [Pei16, Se. 2.3.2], with a limited error growth. Note that there exists a trivial converse reduction from primal-RLWE to dual-RLWE.
- We show that the reduction mentioned above can be extended to a reduction from decision (resp. search) primal-RLWE in K to decision (resp. search) PLWE for f, where K is the field generated by the polynomial f. The analysis is significantly more involved. It requires the introduction of the so-called conductor ideal, to handle the transformation from the ideal \mathcal{O}_K to the order $\mathbb{Z}[x]/f$, and upper bounds on the condition number of the map that sends the coefficient embeddings to the Minkowski embeddings, to show that the noise increases are limited. Our conditioning upper bound is polynomial in n only for limited (but still huge) classes of polynomials that include those of the form $x^n + x \cdot P(x) a$, with deg P < n/2 and a prime that is $\geq 25 \cdot ||P||_1^2$ and $\leq \mathsf{poly}(n)$. A trivial converse reduction goes through for the same f's.
- We exploit the recent technique from [PRSD17] to obtain a search to decision reduction for dual-RLWE.

Concretely, the error rate increases are polynomial in $n = \deg K$, the root discriminant $|\Delta_K|^{1/n}$ and, for the reduction to PLWE, in the root algebraic norm $\mathcal{N}(\mathcal{C}_{\mathbb{Z}[\alpha]})^{1/n}$ of the conductor ideal $\mathcal{C}_{\mathbb{Z}[\alpha]}$ of $\mathbb{Z}[\alpha]$, where α is a root of f defining K. We note that in many cases of interest, all these quantities are polynomially bounded in n. To enjoy these limited error rate growths, the first two reductions require knowledge of specific data related to K, namely, a short element (with respect to the Minkowski embeddings) in the different ideal $(\mathcal{O}_K^{\vee})^{-1}$ and a short element in $\mathcal{C}_{\mathbb{Z}[\alpha]}$. In general, these are hard to compute.

3.2.1 Techniques

The first reduction is derived from [LPR10, Le. 2.15] and [Pei16, Se. 2.3.2]: if it satisfies some arithmetic properties, a multiplication by an element $t \in \mathcal{O}_K$ induces an \mathcal{O}_K -module isomorphism from $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ to $\mathcal{O}_K/q\mathcal{O}_K$. For the reduction to be meaningful, we need t to have small Minkowski embeddings. We prove the existence of such a small t satisfying the appropriate arithmetic conditions, by generalizing the inclusion-exclusion technique developed in [SS13] to study the key generation algorithm of the NTRU signature scheme [HHPW10].

The Lyubashevsky *et al.* bijection works with \mathcal{O}_K^{\vee} and \mathcal{O}_K replaced by arbitrary ideals of K, but this does not provide a bijection from $\mathcal{O}_K/q\mathcal{O}_K$ to $\mathbb{Z}[\alpha]/q\mathbb{Z}[\alpha]$, as $\mathbb{Z}[\alpha]$ may only be an order of \mathcal{O}_K (and not necessarily an ideal). We circumvent this difficulty by using the conductor ideal of $\mathbb{Z}[\alpha]$. Intuitively, the conductor ideal describes the relationship between \mathcal{O}_K and $\mathbb{Z}[\alpha]$. As far as we are aware, this is the first time the conductor ideal is used in the RLWE context. This bijection and the existence of an appropriate multiplier t as above provide a (non-uniform) reduction from primal-RLWE to a variant of PLWE for which the noise terms have small Minkowski embeddings (instead of small polynomial coefficients).

We show that for many number fields, the linear map between polynomial coefficients and Minkowski embeddings has a condition number that is polynomially bounded in n, *i.e.*, the map has bounded distortion and behaves not too noticeably differently from a scaling. This implies that the latter reduction is also a reduction from primal-RLWE to standard PLWE for these rings. We were able to show condition number bounds that are polynomial in n only for restricted families of polynomials f, yet exponentially large as n increases. These include in particular those of the form mentioned above. Note that the primality condition on the constant coefficient is used only to ensure that f is irreducible and hence defines a number field. For these f's, we use Rouché's theorem to prove that the roots are close to the scaled n-th roots of unity $(a^{1/n} \cdot \alpha_n^k)_{0 \le k < n}$, and then that f "behaves" as $x^n - a$ in terms of geometric distortion.

Our search-to-decision reduction for dual-RLWE relies on techniques developed in [PRSD17]. In that article, Peikert et al. consider the following 'oracle hidden center' problem (OHCP). In this problem, we are given access to an oracle \mathcal{O} taking as inputs a vector $\mathbf{z} \in \mathbb{R}^k$ and a scalar $t \in \mathbb{R}^{\geq 0}$, and outputting a bit. The probability that the oracle outputs 1 (over its internal randomness) is assumed to depend only on $\exp(t) \cdot \|\mathbf{z} - \mathbf{x}\|$, for some vector \mathbf{x} . The goal is to recover \mathcal{O} 's center \mathbf{x} . On the one hand, Peikert *et* al. give a polynomial-time algorithm for this problem, assuming the oracle is 'well-behaved' ([PRSD17, Prop. 4.4]). On the other hand, they show how to map a Bounded Distance Decoding (BDD) instance to such an OHCP instance if they have access to Gaussian samples in the dual of the BDD lattice, where the engine of the oracle is the decision dual-RLWE oracle ([PRSD17, Se. 6.1]). We construct the OHCP instance from the decision RLWE oracle in a different manner. We use our input search dual-RLWE samples and take small Gaussian combinations of them. By re-randomizing the secret and adding some noise, we can obtain arbitrarily many dual-RLWE samples. Subtracting from the input samples well-chosen z_i 's in $K_{\mathbb{R}}$ and setting the standard deviation of the Gaussian combination appropriately leads to a valid OHCP instance. The main technical hurdle is to show that a Gaussian combination of elements of $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ is close to uniform. For this, we generalize a ring Leftover Hash Lemma proved for specific pairs (\mathcal{O}_K, q) in [SS11].

3.2.2 Related works

The reductions studied in this chapter can be combined with those from ApproxSVP for \mathcal{O}_K -ideals to dual-RLWE [LPR10, PRSD17]. Recently, Albrecht and Deo [AD17] built upon [BLP⁺13] to obtain a reduction from Module-LWE to RLWE. This can be both combined with our reductions and the quantum reductions from ApproxSVP for \mathcal{O}_K -modules to Module-LWE¹ [LS15, PRSD17]. Downstream, the reductions can be combined with the reduction from PLWE to Middle-Product LWE from Chapter 4. The latter involves an error rate growth that is linearly bounded by the so-called *expansion factor* of f: it turns out that those f's for which we could bound the condition number of the Minkowski map by a polynomial function of deg f also have polynomially bounded expansion factor. These reductions and those considered in this chapter are pictorially described in Figure 3.1.

The ideal-changing scaling element t and the distortion of the Minkowski map were closely studied in [CIV16a, CIV16b, Pei16] for a few precise polynomials and fields. We use the same objects, but provide bounds that work for all (or many) fields.

¹The reduction from [LS15] is limited to cyclotomic fields, but [PRSD17] readily extends to module lattices.



Fig. 3.1: Relationships between variants of RLWE and PLWE. The dotted box contains the problems studied in this chapter. Each arrow may hide a noise rate degradation (and module rank - modulus magnitude transfer in the case of [AD17]). The top to bottom arrows in the dotted box correspond to non-uniform reductions. The reductions involving PLWE are analyzed for limited family of defining polynomials. The arrows without references correspond to trivial reductions.

3.2.3 Impact

As it is standard for the hardness foundations of lattice-based cryptography, our reductions *should not* be considered for setting practical parameters. They should rather be viewed as a strong evidence that the six problems under scope are essentially equivalent and do not suffer from a design flaw (unless they all do). We hope they will prove useful towards understanding the plausibility of weak fields for RLWE.

Our first result shows that there exists a way of reducing dual-RLWE to primal-RLWE while controlling the noise growth. Even though the reduction is non-uniform, it gives evidence that these problems are qualitatively equivalent. Our second result shows that RLWE and PLWE are essentially equivalent for a large class of polynomials/fields. In particular, the transformation map between the Minkowski embeddings and the coefficient embeddings has a bounded distortion. Finally, our search to decision fills an important gap. On the one hand, it precludes the possibility that search RLWE could be harder than decision RLWE. On the other hand, it gives further evidence of the decision RLWE hardness. In [PRSD17], the authors give a reduction from ApproxSVP for \mathcal{O}_K -ideals to decision RLWE. But in the current state of affairs, ApproxSVP for this special class of lattices seems easier than RLWE, at least for some parameters [CDPR16, CDW17], etc. On the opposite, RLWE is qualitatively equivalent to ApproxSIVP for \mathcal{O}_K -modules ([LS15, AD17]).

As the studied problems reduce to one another, one may then wonder which one to use for crypto-

graphic design. Using dual-RLWE requires knowledge of \mathcal{O}_K , which is notoriously hard to compute for an arbitrary field K. This may look as an incentive to use the corresponding PLWE problem instead, as it does not require the knowledge of \mathcal{O}_K . Yet, for it to be useful in cryptographic design, one must be able to decode the noise from its representative modulo a scaled version of the lattice corresponding to $\mathbb{Z}[\alpha]$. This seems to require the knowledge of a good basis of that lattice, which may not be easy to obtain either, depending on the considered polynomial f.

3.2.4 Follow-up work

We now briefly describe the connection between our contribution and [BBPS19], a follow-up work presented at Asiacrypt 2019. In this work, Bolboceanu *et al.* introduced a generalization of the RLWE problem called Order-LWE, where the ambient ring is not the ring of integers of a number field anymore, but rather a full-rank sub-ring of it (*i.e.* an *order*). They showed that Order-LWE is at least as hard as worst-case lattice problems in invertible ideal lattices of the respective order. The PLWE problem associated to the polynomial f is an instance of Order-LWE for the particular order $\mathcal{O} = \mathbb{Z}[x]/f$. As a consequence, their work implies a worst-case hardness result for PLWE different from the one presented in this chapter, essentially by relating it to a different class of lattices than those considered in [LPR10, PRSD17].

3.3 Orders in number fields

An order \mathcal{O} in K is a number ring which is a finite index subring of \mathcal{O}_K . In particular, the ring of integers \mathcal{O}_K is the maximal order in K. Number rings such as $\mathbb{Z}[\alpha]$, with α a root of a defining polynomial f, are of particular interest. In this chapter, we will work with only these two previously mentioned orders.

The conductor of an order \mathcal{O} is defined as the set $\mathcal{C}_{\mathcal{O}} = \{x \in K : x\mathcal{O}_K \subseteq \mathcal{O}\}$. It is contained in \mathcal{O} , and it is both an \mathcal{O} -ideal and an \mathcal{O}_K -ideal: it is in fact the largest ideal with this property. It is never empty, as it contains the index $[\mathcal{O}_K : \mathcal{O}]$. If it is coprime with the conductor, an ideal in \mathcal{O}_K can be naturally considered as an ideal in \mathcal{O} , and reciprocally. This is made precise in the following lemma.

Lemma 3.1 ([Cona, Th. 3.8]). Let \mathcal{O} be an order in K.

- 1. Let I be an \mathcal{O}_K -ideal coprime to $\mathcal{C}_{\mathcal{O}}$. Then $I \cap \mathcal{O}$ is an \mathcal{O} -ideal coprime to $\mathcal{C}_{\mathcal{O}}$ and the natural map $\mathcal{O}/I \cap \mathcal{O} \longrightarrow \mathcal{O}_K/I$ is a ring isomorphism.
- 2. Let J be an \mathcal{O} -ideal coprime to $\mathcal{C}_{\mathcal{O}}$. Then $J\mathcal{O}_K$ is an \mathcal{O}_K -ideal coprime to $\mathcal{C}_{\mathcal{O}}$ and the natural map $\mathcal{O}/J \longrightarrow \mathcal{O}_K/J\mathcal{O}_K$ is a ring isomorphism.
- 3. The set of \mathcal{O}_K -ideals coprime to $\mathcal{C}_{\mathcal{O}}$ and the set of \mathcal{O} -ideals coprime to $\mathcal{C}_{\mathcal{O}}$ are in multiplicative bijection by $I \longmapsto I \cap \mathcal{O}$ and $J \longmapsto J \mathcal{O}_K$.

The above description does not tell how to "invert" the isomorphisms. This can be done by a combination of the following lemmas and passing through the conductor, as we will show later.

Lemma 3.2. Let \mathcal{O} be an order in K and I an \mathcal{O}_K -ideal coprime to the conductor $\mathcal{C}_{\mathcal{O}}$. Then the inclusions $\mathcal{C}_{\mathcal{O}} \subseteq \mathcal{O}$ and $\mathcal{C}_{\mathcal{O}} \subseteq \mathcal{O}_K$ induce isomorphisms $\mathcal{C}_{\mathcal{O}}/I \cap \mathcal{C}_{\mathcal{O}} \simeq \mathcal{O}/I \cap \mathcal{O}$ and $\mathcal{C}_{\mathcal{O}}/I \cap \mathcal{C}_{\mathcal{O}} \simeq \mathcal{O}_K/I$.

Proof. By assumption we have $C_{\mathcal{O}} + I = \mathcal{O}_K$, so that the homomorphism $C_{\mathcal{O}} \to \mathcal{O}_K/I$ is surjective. By Lemma 3.1, the set $I \cap \mathcal{O}$ is an \mathcal{O} -ideal coprime to $C_{\mathcal{O}}$ so that $C_{\mathcal{O}} + I \cap \mathcal{O} = \mathcal{O}$. This implies that the homomorphism $\mathcal{C}_{\mathcal{O}} \to \mathcal{O}/I \cap \mathcal{O}$ is surjective too. Both homomorphisms have kernel $I \cap \mathcal{C}_{\mathcal{O}}$.

Lemma 3.3 ([Cona, Cor. 3.10]). Let \mathcal{O} be an order in K and $\beta \in \mathcal{O}$ such that $\beta \mathcal{O}_K$ is coprime to $\mathcal{C}_{\mathcal{O}}$. Then $\beta \mathcal{O}_K \cap \mathcal{O} = \beta \mathcal{O}$.

3.4 Formal definitions of dual-RLWE, primal-RLWE and PLWE

We now formally define the computational problems that we will study in this chapter.

Definition 3.1 (RLWE and PLWE distributions). Let K a degree n number field defined by f, \mathcal{O}_K its ring of integers, $\Sigma \succ 0$ and $q \geq 2$.

For $s \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$, we define the dual-RLWE distribution $\mathsf{A}_{q,\Sigma}^{\vee}(s)$ as the distribution over $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$ obtained by sampling $a \leftarrow U(\mathcal{O}_K/q\mathcal{O}_K)$, $e \leftarrow D_{\Sigma}^H$ and returning the pair $(a, a \cdot s + e)$.

For $s \in \mathcal{O}_K/q\mathcal{O}_K$, we define the primal-RLWE distribution $\mathsf{A}_{q,\Sigma}(s)$ as the distribution over $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K$ obtained by sampling $a \leftrightarrow U(\mathcal{O}_K/q\mathcal{O}_K)$, $e \leftarrow D_{\Sigma}^H$ and returning the pair $(a, a \cdot s + e)$.

For $s \in \mathbb{Z}_q[x]/f$, we define the PLWE distribution $\mathsf{P}_{q,\Sigma}(s)$ as the distribution over $\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$ obtained by sampling $a \leftrightarrow U(\mathbb{Z}_q[x]/f)$, $e \leftrightarrow D_{\Sigma}$ and returning the pair $(a, a \cdot s + e)$ (with $\mathbb{R}_q = \mathbb{R}/q\mathbb{Z}$).

In the definition above, we identified the support H of D_{Σ}^{H} with $K_{\mathbb{R}}$, and the support \mathbb{R}^{n} of D_{Σ} with $\mathbb{R}[x]/f$. Note that sampling from $\mathsf{A}_{q,\Sigma}^{\vee}(s)$ and $\mathsf{A}_{q,\Sigma}(s)$ seems to require the knowledge of a basis of \mathcal{O}_{K} . It is not known to be computable in polynomial-time from a defining polynomial f of an arbitrary K. In this chapter, we assume that a basis of \mathcal{O}_{K} is known.

Definition 3.2 (The RLWE and PLWE problems). We use the same notations as above. Further, we let \mathcal{E}_{\succ} be a subset of $\Sigma \succ 0$ and D_{\succ} be a distribution over $\Sigma \succ 0$.

Search dual-RLWE_{q, \mathcal{E}_{\succ}} (resp. primal-RLWE and PLWE) consists in finding s from a sampler from $\mathsf{A}_{q,\Sigma}^{\vee}(s)$ (resp. $\mathsf{A}_{q,\Sigma}(s)$ and $\mathsf{P}_{q,\Sigma}(s)$), where $s \in \mathcal{O}_{K}^{\vee}/q\mathcal{O}_{K}^{\vee}$ (resp. $s \in \mathcal{O}_{K}/q\mathcal{O}_{K}$ and $s \in \mathbb{Z}_{q}[x]/f$) and $\Sigma \in \mathcal{E}_{\succ}$ are arbitrary.

Decision dual-RLWE_{q,D} (resp. primal-RLWE and PLWE) consists in distinguishing between a sampler from $A_{q,\Sigma}^{\vee}(s)$ (resp. $A_{q,\Sigma}(s)$ and $P_{q,\Sigma}(s)$) and a uniform sampler over $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$ (resp. $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K$ and $\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$), with non-negligible probability over $s \leftrightarrow \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ (resp. $s \in \mathcal{O}_K/q\mathcal{O}_K$ and $s \in \mathbb{Z}_q[x]/f$) and $\Sigma \leftrightarrow D_{\succ}$.

In Chapter 5, we are going to use a variant of the decision PLWE problem where the error is either drawn from a continuous distribution as above and then rounded to the nearest integer, either drawn from a distribution on $\mathbb{Z}_q[x]/f$. We call these variants *discrete* in contrast with the one defined above which we call *continuous*.

When the distribution D_{\succ} over $\Sigma \succeq 0$ assigns a non-zero probability only to a single positive definite matrix Σ , Decision dual-RLWE_{q,D_{\succ}} (resp. primal-RLWE and PLWE) asks to distinguish between a sampler from $A_{q,\Sigma}^{\vee}(s)$ (resp. $A_{q,\Sigma}(s)$ and $P_{q,\Sigma}(s)$) and a uniform sampler over $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$ (resp. $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K$ and $\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$), with non-negligible probability over $s \leftrightarrow \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ (resp. $s \in \mathcal{O}_K/q\mathcal{O}_K$ and $s \in \mathbb{Z}_q[x]/f$). In this case, we write Decision dual-RLWE_{q,D_Σ} (resp. primal-RLWE and PLWE) instead of Decision dual-RLWE_{q,D_{\succ}} (resp. primal-RLWE and PLWE).

The dual-RLWE (resp. primal-RLWE and PLWE) problems have also been studied for different secret distributions on $\mathcal{O}_{K}^{\vee}/q\mathcal{O}_{K}^{\vee}$ (resp. $\mathcal{O}_{K}/q\mathcal{O}_{K}$ and $\mathbb{Z}_{q}[x]/f$) in [ACPS09, BBPS19], etc. In Chapters 4 and 5, we are going to use a variant of the decision PLWE problem where the error is drawn from a χ_{1} distribution and the secret is drawn from a distribution χ_{2} on $\mathbb{Z}_{q}[x]/f$ that both produce small elements with high ptobability. We use the notation $\mathsf{PLWE}_{q,\chi_{1},\chi_{2}}$ for this variant. It was observed in [LPR10] that under a condition on q which ensures that a uniform element in $\mathbb{Z}_{q}/(f)$ is invertible with non-negligible probability, the reduction from uniform secret to small secret described in [ACPS09] in the context of LWE also applies to PLWE.

Lemma 3.4. Let f be a polynomial of degree n and $q \ge n$ such that the factors of f modulo q are distinct. Let χ_1 and χ_2 be distributions over $\mathbb{Z}_q[x]/(f)$. Then there is a ppt reduction from $\mathsf{PLWE}_{q,\chi_1,\chi_2}^{(f)}$ to $\mathsf{PLWE}_{q,\chi_1,\chi_1}^{(f)}$.

The hardness of PLWE was investigated in [SSTX09, LPR10], among others. The problems above are in fact defined for sequences of number fields of growing degrees n such that the bit-size of the problem description grows at most polynomially in n. The run-times, success probabilities and distinguishing advantages of the algorithms solving the problems are considered asymptotically as functions of n.

We will consider variants of the decision problems for which the distinguishing must occur for all $s \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ (resp. $s \in \mathcal{O}_K/q\mathcal{O}_K$ and $s \in \mathbb{Z}_q[x]/f$) and all $\Sigma \in \mathcal{E}_{\succ}$ rather than with non-negligible probability over s. We call this variant worst-case decision dual-RLWE (resp. primal-RLWE and PLWE). When the set \mathcal{E}_{\succ} consists of a single matrix Σ , we write Decision dual-RLWE_{q,D_{Σ}} (resp. primal-RLWE and PLWE) instead of Decision dual-RLWE_{q,\mathcal{E}_{\succ}} (resp. primal-RLWE and PLWE). Under some conditions on D_{\succ} and \mathcal{E}_{\succ} , these variants are computationally equivalent.

Lemma 3.5 (Adapted from [LPR10, Se. 5.2]). We use the same notations as above. If $\Pr_{\Sigma \leftrightarrow D_{\succ}}[\Sigma \notin \mathcal{E}_{\succ}] \leq 2^{-n}$, then decision dual-RLWE_{q,D_{\sigma}} (resp. primal-RLWE and PLWE) reduces to worst-case decision dual-RLWE_{q,\mathcal{E}_{\sigma}} (resp. primal-RLWE and PLWE).

Assume further that D_{\succ} can be sampled from in polynomial-time. If $\max_{\Sigma \in \mathcal{E}_{\succ}} R(D_{\succ} || D_{\succ} + \Sigma) \leq$ poly(n), then worst-case decision dual-RLWE_{q, $\mathcal{E}_{\succ}} (resp. primal-RLWE and PLWE) reduces to decision dual-RLWE_{q,<math>\mathcal{D}_{\succ}$} (resp. primal-RLWE and PLWE).</sub>

Note that it is permissible to use the Rényi divergence here even though we are considering decision problems. Indeed, the argument is applied to the random choice of the noise distribution and not to the distinguishing advantage. The same argument has been previously used in [LPR10, Se. 5.2].

Proof. The first statement is direct. We prove the second statement only for dual-RLWE, as the proofs for primal-RLWE and PLWE are direct adaptations. Assume we are given a sampler that outputs (a_i, b_i) with $a_i \leftrightarrow U(\mathcal{O}_K/q\mathcal{O}_K)$ and b_i either uniform in $K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$ or of the form $b_i = a_i s + e_i$ with $s \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ and $e_i \leftrightarrow D_{\Sigma}^H$. The reduction proceeds by sampling $s' \leftrightarrow U(\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee})$ and $\Sigma' \leftrightarrow D_{\succ}$, and mapping all input (a_i, b_i) 's to $(a'_i, b'_i) = (a_i, b_i + a_i s' + e'_i)$ with $e'_i \leftrightarrow D_{\Sigma'}^H$. This transformation maps the uniform distribution to itself, and $\mathsf{A}_{q,\Sigma}^{\vee}(s)$ to $\mathsf{A}_{q,\Sigma''}^{\vee}(s+s')$ with $\Sigma''_{ij} = \Sigma_{ij} + \Sigma'_{ij}$ for all i, j. If the success probability (success being enjoying a non-negligible distinguishing advantage) over the error parameter sampled from D_{\succ} is non-negligible, then so is it for the error parameter sampled $D_{\succ} + \Sigma$, as, by assumption, the Rényi divergence $R(D_{\succ}||D_{\succ} + \Sigma)$ is polynomially bounded.

Many choices of D_{\succ} and \mathcal{E}_{\succ} satisfy the conditions of Lemma 3.5. The following is inspired from [LPR10, Se. 5.2]. We define the distribution \mathcal{E}_{\succ} as follows, for an arbitrary r: Let $s_{ij} = r^2(1+nx_{ij})$ for all i > j, $s_{ii} = r^2(1+n^3x_{ii})$ for all i and $s_{ij} = s_{ji}$ for all i < j, where the x_{ij} 's are independent samples from the $\Gamma(2, 1)$ distribution (of density function $x \mapsto x \exp(-x)$); the output matrix is $(s_{ij})_{ij}$. Note that it is symmetric and strictly diagonally dominant (and hence $\succ 0$) with probability $1-2^{-\Omega(n)}$. Then the set of all $\Sigma \succ 0$ with coefficients of magnitudes $\leq r^2 n^4$ satisfies the first condition of Lemma 3.5, and the set of all $\Sigma \succ 0$ with coefficients of magnitudes $\leq r^2$ satisfies the second condition of Lemma 3.5. We can hence switch from one variant to the other while incurring an error rate increase that is $\leq \operatorname{poly}(n)$.

3.5 From dual-RLWE to primal-RLWE

The following result is the key ingredient for the dual-RLWE to primal-RLWE and primal-RLWE to PLWE reductions.

Lemma 3.6 ([LPR10, Le. 2.14]). Let I and J two \mathcal{O}_K -ideals. Let $t \in I$ such that the ideals $t \cdot I^{-1}$ and J are coprime and let \mathcal{M} be any fractional \mathcal{O}_K -ideal. Then the function $\theta_t : \mathcal{M} \to \mathcal{M}$ defined as $\theta_t(x) = t \cdot x$ induces an \mathcal{O}_K -module isomorphism from $\mathcal{M}/J\mathcal{M}$ to $I\mathcal{M}/IJ\mathcal{M}$. The authors of [LPR10] also gave an explicit way to obtain a suitable t by solving a set of conditions stemming from the Chinese Remainder Theorem. However, this construction does not give good control on the magnitudes of the Minkowski embeddings of t. In Chapter 3 we show that the size of t can be controlled by Gaussian sampling.

Theorem 3.1 (Adapted from [Pei16, Se. 2.3.2]). Let $\Sigma \succ 0$ and $s \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$. Let $t \in (\mathcal{O}_K^{\vee})^{-1}$ such that $t(\mathcal{O}_K^{\vee}) + q\mathcal{O}_K = \mathcal{O}_K$. Then the map $(a, b) \mapsto (a, t \cdot b)$ transforms $\mathsf{A}_{q,\Sigma}^{\vee}(s)$ to $\mathsf{A}_{q,\Sigma'}(t \cdot s)$ and $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee})$ into $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K)$, with $\Sigma' = \operatorname{diag}(|\sigma_i(t)|) \cdot \Sigma \cdot \operatorname{diag}(|\sigma_i(t)|)$. The natural inclusion $\mathcal{O}_K \to \mathcal{O}_K^{\vee}$ induces a map that transforms $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K)$ to $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K)$, and $\mathsf{A}_{q,\Sigma}(s)$ to $\mathcal{A}_{a,\Sigma}^{\vee}(s)$.

Proof. First, let $(a, b = a \cdot s + e)$ be distributed as $\mathcal{A}_{q,\Sigma}^{\vee}(s)$. We define $b' = t \cdot b = a \cdot (t \cdot s) + e'$, with $e' = t \cdot e$. By Lemma 3.6, multiplication by t induces an \mathcal{O}_K -module isomorphism $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee} \simeq \mathcal{O}_K/q\mathcal{O}_K$, hence $t \cdot s \in \mathcal{O}_K/q\mathcal{O}_K$. Also, the distribution of the error term e' is $D_{\Sigma'}^H$. As a consequence, the sample (a, b') is distributed as $\mathcal{A}_{q,\Sigma'}(t \cdot s)$. Second, if (a, b) is uniform in $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$, as multiplying by t induces an isomorphism, we have that b' is uniform in $K_{\mathbb{R}}/q\mathcal{O}_K$, independently from a. For the converse reduction, we map $(a, b) \in \mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K$ to $(a, b \mod q\mathcal{O}_K^{\vee}) \in \mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$. Since $q\mathcal{O}_K \subseteq q\mathcal{O}_K^{\vee}$, the map is well defined and it also maps $A_{q,\Sigma}(s)$ to $\mathcal{A}_{q,\Sigma}^{\vee}(s)$ and $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K)$ to $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee})$.

3.6 Controlling the noise growth in the dual to primal reduction

The reduction of Theorem 3.1 is built upon the existence of t as in Lemma 3.6. While this existence is guaranteed constructively by [LPR10], the size is not controlled by the construction. Another t that satisfies the conditions is $t = f'(\alpha)$, where f' is the derivative of f defining $K = \mathbb{Q}[\alpha]$. Indeed, from [Conb, Rem. 4.5], we know that $f'(\alpha) \in (\mathcal{O}_K^{\vee})^{-1}$. However, the noise growth incurred by multiplication by $f'(\alpha)$ may be rather large in general: we have $N(f'(\alpha)) = \Delta_f = [\mathcal{O}_K : \mathbb{Z}[\alpha]]^2 \cdot \mathcal{N}((\mathcal{O}_K^{\vee})^{-1})$.

In this section, we give a probabilistic proof that adequate t's with controlled size can be found by Gaussian sampling.

Let I and J be integral ideals of \mathcal{O}_K . Theorem 3.2 below states that a Gaussian sample t in I is such that $t \cdot I^{-1} + J = \mathcal{O}_K$ with non-negligible probability. The main technical hurdle is to show that the sample is not trapped in IJ' with J' a non-trivial factor of J. We handle this probability in different ways depending on the algebraic norm of J', extending an idea used in [SS13, Se. 4].

- For small-norm factors J' of J, the Gaussian folded modulo IJ' is essentially uniform over I/IJ', by Lemma 2.11. This requires the standard deviation parameter s to be above the smoothing parameter of IJ'. We use the smoothing parameter bound from Lemma 2.14.
- For large-norm factors J', we argue that the non-zero points of IJ' are very unlikely to be hit, thanks to the Gaussian tail bound given in Lemma 2.10 and the fact that the lattice minimum of IJ' is large, by Lemma 2.13.
- For middle-norm factors J', neither of the arguments above applies. Instead, we bound the probability that t belongs to IJ' by the probability that t belongs to IJ'', where J'' is a non-trivial factor of J', and use the first argument above. The factor J'' must be significantly denser than J' so that we have smoothing. But it should also be significantly sparser than \mathcal{O}_K so that the upper bound is not too large.

Setting the standard deviation parameter of the discrete Gaussian so that at least one of the three arguments above applies is non-trivial. In particular, this highly depends on how the ideal J factors

into primes (whether the pieces are numerous, balanced, unbalanced, etc). The choice we make below works in all cases while still providing a reasonably readable proof and still being sufficient for our needs, from an asymptotic perspective. In many cases, better choices can be made. If J is prime, we can take a very small s and use only the second argument. If all factors of J are small, there is good enough 'granularity' in the factorization to use the third argument, and again s can be chosen very small.

Theorem 3.2. Let I and J be integral \mathcal{O}_K -ideals, and write $J = \mathfrak{p}_1^{e_1} \dots \mathfrak{p}_k^{e_k}$ for some prime ideals \mathfrak{p}_i . We sort the \mathfrak{p}_i 's by non-decreasing algebraic norms. Assume that we can take $\delta \in [\frac{4n + \log_2 \Delta_K}{\log_2 \mathcal{N}(J)}, 1]$.² We define:

$$s = \begin{cases} \left(\mathcal{N}(J)^{1/2}\mathcal{N}(I)\Delta_{K}\right)^{1/n} & \text{if} \quad \mathcal{N}(\mathfrak{p}_{k}) \geq \mathcal{N}(J)^{1/2+\delta}, \\ \left(\mathcal{N}(J)^{1/2+2\delta}\mathcal{N}(I)\Delta_{K}\right)^{1/n} & \text{else.} \end{cases}$$

Then we have

$$\Pr_{t \leftarrow D_{I,s}}[tI^{-1} + J = \mathcal{O}_K] \ge 1 - \frac{k}{\mathcal{N}(\mathfrak{p}_1)} - 2^{-n+4}.$$

Proof. We bound the probability P of the negation, from above. We have

$$P = \Pr_{t \leftarrow D_{I,s}} [t \in \bigcup_{i \in [k]} I\mathfrak{p}_i] = \sum_{S \subseteq [k], S \neq \emptyset} (-1)^{|S|+1} \cdot \Pr_{t \leftarrow D_{I,s}} [t \in I \cdot \prod_{i \in S} \mathfrak{p}_i].$$

We rewrite it as $P = P_1 + P_2$ with

$$P_{1} = \sum_{S \subseteq [k], S \neq \emptyset} (-1)^{|S|+1} \frac{1}{\prod_{i \in S} \mathcal{N}(\mathfrak{p}_{i})} = 1 - \prod_{i \in [k]} \left(1 - \frac{1}{\mathcal{N}(\mathfrak{p}_{i})}\right),$$
$$P_{2} = \sum_{S \subseteq [k], S \neq \emptyset} (-1)^{|S|+1} \left(\Pr_{t \leftrightarrow D_{I,s}}[t \in I \cdot \prod_{i \in S} \mathfrak{p}_{i}] - \prod_{i \in S} \frac{1}{\mathcal{N}(\mathfrak{p}_{i})}\right).$$

We have $P_1 \leq 1 - (1 - 1/\mathcal{N}(\mathfrak{p}_1))^k \leq k/\mathcal{N}(\mathfrak{p}_1)$. Our task is now to bound P_2 .

Assume first that $\mathcal{N}(\mathfrak{p}_k) \geq \mathcal{N}(J)^{1/2+\delta}$. This implies that $\prod_{i \in S} \mathcal{N}(\mathfrak{p}_i) \leq \mathcal{N}(J)^{1/2-\delta}$ for all $S \subseteq [k]$ not containing k. By Lemma 2.14, we have $s \geq \eta_{\epsilon}(I \prod_{i \in S} \mathfrak{p}_i)$ for all such S's, with $\epsilon = 2^{-2n}$. We "smooth" out those ideals, *i.e.*, we use Lemma 2.11 to obtain, for all $S \subseteq [k] \setminus \{k\}$:

$$\left| \Pr_{t \leftrightarrow D_{I,s}} [t \in I \cdot \prod_{i \in S} \mathfrak{p}_i] - \prod_{i \in S} \frac{1}{\mathcal{N}(\mathfrak{p}_i)} \right| \le 2\epsilon.$$

Now if S is a subset containing k, then we have $\mathcal{N}(\prod_{i\in S}\mathfrak{p}_i) \geq \mathcal{N}(J)^{1/2+\delta}$. By Lemma 2.13, we have $\lambda_1(I\prod_{i\in S}\mathfrak{p}_i) \geq \sqrt{n}\cdot\mathcal{N}(I)^{1/n}\mathcal{N}(J)^{(1/2+\delta)/n}$. On the other hand, by Lemma 2.10, we have $\Pr_{t\leftrightarrow D_{I,s}}[||t|| \geq 2s\sqrt{n}] \leq 2^{-2n}$. Thanks to our choice of s, the assumption on δ and Lemma 2.12, we obtain

$$\Pr_{t \leftarrow D_{I,s}}[t \in I \prod_{i \in S} \mathfrak{p}_i] \le \Pr_{t \leftarrow D_{I,s}}[t=0] + 2^{-2n} \le 2^{-2n+2}.$$

This allows us to bound P_2 as follows:

$$P_2 \le 2^k \cdot \left(\epsilon + 2^{-2n+2} + \mathcal{N}(J)^{-(1/2+\delta)}\right).$$

By assumption on δ , we have $\mathcal{N}(J) \geq 2^{2n}$ and $P_2 \leq 2^{-n+3}$. This completes the proof for the large $\mathcal{N}(\mathfrak{p}_k)$ case.

²The parameter δ should be thought as near 0. It can actually be chosen such if $\mathcal{N}(J)$ is sufficiently large.

Now, assume that $\mathcal{N}(\mathfrak{p}_k) < \mathcal{N}(J)^{1/2+2\delta}$. Then, as above, the definition of s implies that, for any $S \subseteq [k]$ with $\mathcal{N}(\prod_{i \in S} \mathfrak{p}_i) \leq \mathcal{N}(J)^{1/2+\delta}$, we have $|\Pr[t \in I \prod_{i \in S} \mathfrak{p}_i] - 1/\prod_{i \in S} \mathcal{N}(\mathfrak{p}_i)| \leq 2^{-2n+1}$. Also as above, if we have $\mathcal{N}(\prod_{i \in S} \mathfrak{p}_i) \geq \mathcal{N}(J)^{1/2+3\delta}$, then $\lambda_1(I \prod_{i \in S} \mathfrak{p}_i)$ is too large for a non-zero element of $I \prod_{i \in S} \mathfrak{p}_i$ to be hit with significant probability. Assume finally that

$$\mathcal{N}(J)^{1/2+2\delta} \leq \mathcal{N}(\prod_{i \in S} \mathfrak{p}_i) \leq \mathcal{N}(J)^{1/2+3\delta}.$$

As $\mathcal{N}(\mathfrak{p}_k) < \mathcal{N}(J)^{1/2+\delta}$, there exists $S' \subseteq S$ such that

$$\mathcal{N}(J)^{\delta} \leq \mathcal{N}(\prod_{i \in S'} \mathfrak{p}_i) \leq \mathcal{N}(J)^{1/2+2\delta}$$

By inclusion, we have that $\Pr[t \in I \prod_{i \in S} \mathfrak{p}_i] \leq \Pr[t \in I \prod_{i \in S'} \mathfrak{p}_i]$. Now, as the norm of $\prod_{i \in S'} \mathfrak{p}_i$ is small enough, we can use the smoothing argument above to claim that

$$\Pr_{t \leftarrow D_{I,s}}[t \in I \prod_{i \in S'} \mathfrak{p}_i] \le 2^{-2n+1} + \frac{1}{\mathcal{N}(\prod_{i \in S'} \mathfrak{p}_i)} \le 2^{-2n+1} + \frac{1}{\mathcal{N}(J)^{\delta}}.$$

By assumption on δ , the latter is $\leq 2^{-n+2}$. Collecting terms allows to complete the proof.

The next corollary shows that the needed t can be found with non-negligible probability.

Corollary 3.1. Let I be an integral \mathcal{O}_K -ideal. Let $q \ge \max(2n, 2^{16} \cdot \Delta_K^{8/n})$ be a prime rational integer and \mathfrak{p}_k a prime factor of $q\mathcal{O}_K$ with largest norm. We define:

$$s = \begin{cases} q^{1/2} \cdot (\mathcal{N}(I)\Delta_K)^{1/n} & \text{if} \quad \mathcal{N}(\mathfrak{p}_k) \ge q^{(5/8) \cdot n}, \\ q^{3/4} \cdot (\mathcal{N}(I)\Delta_K)^{1/n} & \text{else.} \end{cases}$$

Then, for sufficiently large n, we have

$$\Pr_{t \leftarrow D_{I,s}}[tI^{-1} + q\mathcal{O}_K = \mathcal{O}_K] \ge 1/2.$$

Proof. The result follows from applying Theorem 3.2 with $J = q\mathcal{O}_K$ and $\delta = 1/8$. The first lower bound on q ensures that $k/\mathcal{N}(\mathfrak{p}_1) \leq 1/2$, where $k \leq n$ denotes the number of prime factors of $q\mathcal{O}_K$ and \mathfrak{p}_1 denotes a factor with smallest algebraic norm. The second lower bound on q ensures that we can indeed set $\delta = 1/8$.

We insist again on the fact that the required lower bounds on s can be much improved under specific assumptions on the factorization of q. For example, one could choose a q such that all the factors of $q\mathcal{O}_K$ have large norms, by sampling q randomly and checking its primality and the factorization of the defining polynomial f modulo q. In that case, the factors $q^{1/2}$ and $q^{3/4}$ can be decreased drastically.

We note that if the noise increase incurred by a reduction from an LWE-type problem to another is bounded as $n^{c_1} \cdot q^{c_2}$ for some $c_1 < 1$ and some $c_2 < 1$, then one may set the working modulus q so that the starting LWE problem has a sufficient amount of noise to not be trivially easy to solve, and the ending LWE problem has not enough noise to be information-theoretically impossible to solve (else the reduction would be vacuous). Indeed, it suffices to set q sufficiently larger than $n^{c_1/(1-c_2)}$.

3.7 From primal-RLWE to PLWE

In this section, we describe a reduction from primal-RLWE to PLWE. As an intermediate step, we first consider a reduction from primal-RLWE to a variant $PLWE^{\sigma}$ of PLWE where the noise is small with

respect to the Minkowski embedding rather than the coefficient embedding. Then, we assess the noise distortion when looking at its Minkowski embedding versus its coefficient embedding.

If $K = \mathbb{Q}[x]/f$ for some $f = \prod_{j \leq n} (x - \alpha_j)$, the associated Vandermonde matrix V_f has *j*th row $(1, \alpha_j, \ldots, \alpha_j^{n-1})$ and corresponds to the linear map between the coefficient and Minkowski embedding spaces. Thus a good approximation of the distortion is given by the condition number $\text{Cond}(V_f) = s_n/s_1$, where the s_i 's refer to the largest/smallest singular values of V_f .

Since $\operatorname{Cond}(V_f) \leq ||V_f||_{\mathsf{F}} \cdot ||V_f^{-1}||_{\mathsf{F}}$, these matrix norms also quantify how much V_f distorts the space. For a restricted, yet exponentially large, family of polynomials defining number fields, we show that both $||V_f||_{\mathsf{F}}$ and $||V_f^{-1}||_{\mathsf{F}}$ are polynomially bounded.

To do this, we start from $f_{n,a} = x^n - a$ whose distortion is easily computable. Then we add a "small perturbation" to this polynomial. Intuitively, the roots of the resulting polynomial should not move much, so that the norms of the "perturbed" Vandermonde matrices should be essentially the same. We formalize this intuition in Section 3.7.2 and locate the roots of the perturbed polynomial using Rouché's theorem.

Mapping a sample of PLWE^{σ} to a sample of the corresponding PLWE simply consists in changing the geometry of the noise distribution. A noise distribution with covariance matrix Σ in the Minkowski embedding corresponds to a noise distribution of covariance matrice $(V_f^{-1})^T \Sigma V_f^{-1}$ in the coefficient space. The converse is also true, replacing V_f^{-1} by V_f . Moreover, the noise growths incurred by the reductions remain limited whenever $\|V_f\|_{\mathsf{F}}$ and $\|V_f^{-1}\|_{\mathsf{F}}$ are small.

Overall, reductions between primal-RLWE to PLWE can be obtained by combining Theorems 3.3 and 3.5 below (with Lemma 3.5 to randomize the noise distributions).

3.7.1 Reducing primal-RLWE to $PLWE^{\sigma}$

We keep the notations of the previous section, and let $\mathbb{Z}[x]/(f) = \mathcal{O}$.

Definition 3.3 (The PLWE^{σ} problem). Let also Σ be a positive definite matrix, and $q \geq 2$. For $s \in \mathcal{O}/q\mathcal{O}$, we define the PLWE^{σ} distribution $\mathsf{P}_{q,\Sigma}^{\sigma}(s)$ as the distribution over $\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O}$ obtained by sampling $a \leftarrow U(\mathcal{O}/q\mathcal{O})$, $e \leftarrow D_{\Sigma}^{H}$ and returning the pair $(a, a \cdot s + e)$

Let D_{\succ} be a distribution over $\Sigma \succ 0$. Decision PLWE^{σ} consists in distinguishing between a sampler from $\mathsf{P}_{q,\Sigma}^{\sigma}(s)$ and a uniform sampler over $\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O}$, with non-negligible probability over $s \leftarrow \mathcal{O}/q\mathcal{O}$ and $\Sigma \leftarrow D_{\succ}$.

Theorem 3.3. Assume that $q\mathcal{O}_K + \mathcal{C}_{\mathcal{O}} = \mathcal{O}_K$. Let Σ be a positive definite matrix and $s \in \mathcal{O}_K/q\mathcal{O}_K$. Let $t \in \mathcal{C}_{\mathcal{O}}$ such that $t\mathcal{C}_{\mathcal{O}}^{-1} + q\mathcal{O}_K = \mathcal{O}_K$. Then the map $(a,b) \mapsto (t \cdot a, t^2 \cdot b)$ transforms $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K)$ to $U(\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O})$ and $\mathsf{A}_{q,\Sigma}(s)$ to $\mathsf{P}_{q,\Sigma'}^{\sigma}(t \cdot s)$, where the new covariance is $\Sigma' = \mathsf{diag}(|\sigma(t_i)|^2) \cdot \Sigma \cdot \mathsf{diag}(|\sigma_i(t)|^2)$.

Let $\mathsf{P}_{q,\Sigma}^{\sigma}(s)$ be a PLWE^{σ} distribution. The natural inclusion $\mathcal{O} \to \mathcal{O}_K$ induces a map that transforms $U(\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O})$ to $U(\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K)$ and $\mathsf{P}_{q,\Sigma}^{\sigma}(s)$ to $\mathsf{A}_{q,\Sigma}(s)$.

Proof. Let $(a, b = a \cdot s + e)$ be distributed as $A_{q,\Sigma}(s)$. Let $a' = t \cdot a$ and $b' = t^2 \cdot b = a' \cdot (t \cdot s) + e'$, with $e' = t^2 \cdot e$. Then a' is uniformly distributed in $\mathcal{C}_{\mathcal{O}}/q\mathcal{C}_{\mathcal{O}}$ by applying Lemma 3.6 for $I = \mathcal{C}_{\mathcal{O}}, J = q\mathcal{O}_K$ and $\mathcal{M} = \mathcal{O}_K$. It is also uniformly distributed in $\mathcal{O}/q\mathcal{O}$ by combining Lemma 3.2 and Lemma 3.3. The noise follows the claimed distribution, see the observation in Section 2.4.5. The fact that $t \cdot s \in \mathcal{O}/q\mathcal{O}$ completes the proof that $A_{q,\Sigma}(s)$ is mapped to $\mathsf{P}_{q,\Sigma'}^{\sigma}(t \cdot s)$.

Now, let (a, b) be uniform in $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K$. We already know that a' is uniformly distributed in $\mathcal{O}/q\mathcal{O}$. Let us now consider the distribution of b'. Thanks to the assumption on $q\mathcal{O}_K$, we also have $t^2\mathcal{C}_{\mathcal{O}}^{-1} + q\mathcal{O}_K = \mathcal{O}_K$. Therefore, by Lemma 3.6, multiplication by t^2 induces an isomorphism $\mathcal{O}_K/q\mathcal{O}_K \simeq \mathcal{C}_{\mathcal{O}}/q\mathcal{C}_{\mathcal{O}}$, and hence, by Lemmas 3.2 and 3.3, an isomorphism $\mathcal{O}_K/q\mathcal{O}_K \simeq \mathcal{O}/q\mathcal{O}$. This gives the first reduction. We now turn to the converse reduction. By coprimality and Lemmas 3.2 and 3.6, we have $|\mathcal{O}/q\mathcal{O}| = |\mathcal{O}_K/q\mathcal{O}_K|$. This implies that, thanks to the inclusion $\mathcal{O} \subseteq \mathcal{O}_K$, any (a, b) uniform in $\mathcal{O}/q\mathcal{O} \times K_{\mathbb{R}}/q\mathcal{O}$ is also uniform in $\mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K$ and $\mathsf{P}_{q,\Sigma}^{\sigma}(s)$ is mapped to $\mathsf{A}_{q,\Sigma}(s)$.

As Theorem 3.1, Theorem 3.3 relies on a the existence of a good multiplier. Writing $K = \mathbb{Q}[x]/(f) = \mathbb{Q}[\alpha]$ and $\mathcal{O} = \mathbb{Z}[\alpha]$, the element $f'(\alpha)$ again satisfies the constraints. Indeed, we know that $\mathcal{O}^{\vee} = \frac{1}{f'(\alpha)}\mathcal{O}$ (see [Conb, Th. 3.7]), and we have the inclusion $\mathcal{O}_K \subseteq \mathcal{O}^{\vee}$. Multiplying by $f'(\alpha)$, we obtain $f'(\alpha)\mathcal{O}_K \subseteq \mathcal{O}$. By definition, this means that $f'(\alpha) \in \mathcal{C}_{\mathcal{O}}$, as claimed. While a large $f'(\alpha)$ would mean a large noise growth in the primal-RLWE to PLWE^{σ} reduction, we described in Section 3.6 how to find a smaller adequate multiplier if needed.

We have $\mathcal{N}(f'(\alpha)) = [\mathcal{O}_K : \mathbb{Z}[\alpha]]^2 \cdot \Delta_K$, and, from [Ste17, p.48], the prime factors of $[\mathcal{O}_K : \mathbb{Z}[\alpha]]$ are exactly those of $\mathcal{N}(\mathcal{C}_{\mathcal{O}})$. Provided the valuations are not too high, there should be smaller elements in $\mathcal{C}_{\mathcal{O}}$ than $f'(\alpha)$. We provide in Section 3.8 concrete examples of number fields with defining polynomials fsuch that the norm of $f'(\alpha)$ is considerably larger than both the norms of $\mathcal{C}_{\mathcal{O}}$ and $(\mathcal{O}_K^{\vee})^{-1}$.

3.7.2 Distortion between embeddings

To bound the norms of a Vandermonde matrix associated to a polynomial and its inverse, we study the magnitude of the roots and their pairwise distances. It is known that the Frobenius norm $||V||_{\mathsf{F}}$ of a matrix V satisfies $||V||_{\mathsf{F}}^2 = \operatorname{Tr}(V^*V)$, where * denotes the transpose-conjugate operator. For Vandermonde matrices, this gives

$$\|V_f\|_{\mathsf{F}}^2 = \sum_{j \in [n]} \sum_{k \in [n]} |\alpha_j|^{2(k-1)},$$
(3.1)

which can be handled when the magnitudes of the α_j 's are known. The entries of $V_f^{-1} = (w_{ij})$ have well-known expressions as:

$$w_{ij} = (-1)^{n-i} \frac{e_{n-i}(\overline{\alpha}^j)}{\prod\limits_{k \neq j} (\alpha_j - \alpha_k)},$$
(3.2)

where $e_0 = 1$, e_j for j > 0 stands for the elementary symmetric polynomial of total degree j in n-1 variables, and $\overline{\alpha}^j = (\alpha_1, \ldots, \alpha_{j-1}, \alpha_{j+1}, \ldots, \alpha_n)$, the vector of all roots but α_j . We have the following useful relations with the symmetric functions E_i of all the roots (for all j):

$$\begin{cases} E_1(\boldsymbol{\alpha}) = \alpha_j + e_1(\overline{\boldsymbol{\alpha}}^j), \\ E_i(\boldsymbol{\alpha}) = \alpha_j e_{i-1}(\overline{\boldsymbol{\alpha}}^j) + e_i(\overline{\boldsymbol{\alpha}}^j) \text{ for } 2 \le i \le n-1, \\ E_n(\boldsymbol{\alpha}) = \alpha_j e_{n-1}(\overline{\boldsymbol{\alpha}}^j). \end{cases}$$
(3.3)

Combining (3.3) with Vieta's formulas, bounds on the magnitudes of the roots leads to bounds on the numerators of the w_{ij} 's. The denominators encode the separation of the roots, and deriving a precise lower bound turns out to be the main difficulty. By differentiating $f(x) = \prod_{j \in [n]} (x - \alpha_j)$, we note that $\prod_{k \neq j} |\alpha_j - \alpha_k| = |f'(\alpha_j)|$.

3.7.3 A family of polynomials with easily computable distortion

We first introduce a family of polynomials for which $\|V_f\|_{\mathsf{F}}$ and $\|V_f^{-1}\|_{\mathsf{F}}$ are both simple to estimate. For $n \ge 2$ and $a \ge 1$, we define $f_{n,a} = x^n - a$. The roots can be written³ as $\alpha_j = a^{1/n} e^{2i\pi \frac{j}{n}}$, for $0 \le j < n$.

 $^{^{3}}$ For the rest of this section, 'i' will refer to the imaginary unit.

As these are scalings of the roots of unity, both their magnitude and separation are well-understood. With (3.1), we obtain $\|V_{f_{n,a}}\|_{\mathsf{F}} \leq na^{\frac{n-1}{n}} \leq na$.

For any j, we readily compute $|f'_{n,a}(\alpha_j)| = na^{\frac{n-1}{n}}$. Using (3.3), we observe that $|e_i(\overline{\alpha}^j)| = |\alpha_j|^i$ for $1 \le i < n$. We obtain that the row norm of $V_{f_{n,a}}^{-1}$ is given by its first row as

$$\sum_{j \in [n]} |w_{1j}| = \frac{1}{na^{\frac{n-1}{n}}} \cdot \sum_{j \in [n]} |\alpha_j|^{n-1} = 1,$$

from which it follows that $\|V_{f_{n,a}}^{-1}\|_{\mathsf{F}} \leq \sqrt{n}$.

Next, we show that if we apply a small perturbation to the coefficients of $f_{n,a}$, the norms of the Vandermonde matrix and its inverse are still polynomially bounded. Let $P(x) = \sum_{1 \le j \le \rho \cdot n} p_j x^j$ for some constant $\rho \in (0, 1)$, where the p_j 's are a priori complex numbers. Locating the roots of $g_{n,a} = f_{n,a} + P$ is our first step towards estimating $\|V_{g_{n,a}}\|_{\mathsf{F}}$ and $\|V_{g_{n,a}}^{-1}\|_{\mathsf{F}}$. We will use the following version of Rouché's theorem.

Theorem 3.4 (Rouché, adapted from [Con95, p.125-126]). Let f, P be complex polynomials, and let D be a disk in the complex plane. If for any z on the boundary ∂D we have |P(z)| < |f(z)|, then f and f + P have the same number of zeros inside D, where each zero is counted as many times as its multiplicity.

The lemma below allows to determine sufficient conditions on the parameters such that the assumptions of Theorem 3.4 hold. We consider small disks $D_k = D(\alpha_k, 1/n)$ of radius 1/n around the roots $\alpha_1, \ldots, \alpha_n$ of $f_{n,a}$, and we let ∂D_k denote their respective boundaries. We let $||P||_1 = \sum_j |p_j|$ denote the 1-norm of P.

Lemma 3.7. We have, for all $k \leq n$ and $z \in \partial D_k$:

$$|P(z)| \le (ae)^{\rho} \cdot ||P||_1$$
 and $|f_{n,a}(z)| \ge a \left(1 - \cos(a^{-1/n}) - \frac{2e^{a^{-1/n}}}{na^{2/n}}\right).$

Proof. Write $z = \alpha_k + \frac{e^{it}}{n}$ for some $t \in [0, 2\pi)$. We have $|z| \leq a^{1/n} + 1/n$, and hence $|z|^{\rho n} \leq a^{\rho} \left(1 + \frac{1}{na^{1/n}}\right)^{\rho n}$. The first claim follows from the inequality $|P(z)| \leq \max(1, |z|^{\rho n}) \cdot ||P||_1$.

Next, we have $|f_{n,a}(z)| = a|(1 + \frac{e^{it'}}{na^{1/n}})^n - 1|$, where $t' = t - 2k\pi/n$. W.l.o.g., we assume that k = 0. Let Log denote the complex logarithm, defined on $\mathbb{C} \setminus \mathbb{R}^-$. Since the power series $\sum_{k\geq 1} (-1)^{k-1} u^k/k$ converges to Log(1+u) on the unit disk, we have $\text{Log}(1 + \frac{e^{it}}{na^{1/n}}) = \frac{e^{it}}{na^{1/n}} + \delta$, for some δ satisfying $|\delta| \leq |u| \cdot \sum_{k\geq 1} |u|^k/(k+1) \leq |u|^2$ for $u = \frac{e^{it}}{na^{1/n}}$ (note that it has modulus $\leq 1/n \leq 1/2$). Similarly, we can write $\exp(n\delta) = 1 + \epsilon$ for some ϵ satisfying $|\epsilon| \leq 2n|\delta| \leq 2/(na^{2/n})$. We hence have:

$$|f_{n,a}(z)| = a \cdot |A(t) \cdot (1+\epsilon) - 1| \ge a \cdot ||A(t) - 1| - |\epsilon \cdot A(t)||,$$

with $A(t) = \exp(e^{it}a^{-1/n})$.

Elementary calculus leads to the inequalities $|A(t) - 1| > 1 - \cos(a^{-1/n})$ and $|A(t)| \le e^{a^{-1/n}}$ for all $t \in [0, 2\pi)$. Define $A(t) = \exp(e^{it}a^{-1/n})$ for $t \in [-\pi, \pi]$. We have

$$\arg A(-t) = -\arg A(t) = a^{-1/n} \sin(-t),$$

$$|A(-t)| = |A(t)| = \exp(a^{-1/n} \cos(t)).$$

Therefore, the graph of A(t) is symmetric with respect to the real axis. We can hence restrict the study of A(t) to $[0, \pi]$. As |A(t)| decreases for such t's, this implies that $|A(t)| \le |A(\pi)| \le e^{a^{-1/n}}$ for all t.

Let $\Re A(t)$ and $\Im A(t)$ respectively denote the real and imaginary parts of A(t). Their derivatives are $-\exp(a^{-1/n}\cos(t))a^{-1/n}\cdot\sin(t+a^{-1/n}\sin(t))$ and $\exp(a^{-1/n}\cos(t))a^{-1/n}\cdot\cos(t+a^{-1/n}\sin(t))$, respectively. The study of their signs shows that $\Re A(t)$ decreases on $[0,\pi]$, and that there exists a $t_0 \in (\pi/4, \pi/2)$ such that $\Im A(t)$ increases on $[0, t_0]$ and decreases on $[t_0, \pi]$. We have:

- when $t \in [\pi/2, \pi]$, $\Re A(t) \le \Re A(\pi/2)$ so that $|A(t) 1| \ge 1 \cos(a^{-1/n})$,
- when $t \in [\pi/4, \pi/2]$, $\Im A(t) \ge \min{\{\Im A(\pi/2), \Im A(\pi/4)\}}$ so that

$$|A(t) - 1| \ge \min\{\sin(a^{-1/n}), e^{\sqrt{2}/(2a^{1/n})}\sin(\frac{\sqrt{2}}{2}a^{-1/n})\} \ge 1 - \cos(a^{-1/n}),$$

• when $t \in [0, \pi/4]$, $\Re A(t) \ge \Re A(\pi/4)$, so that

$$|A(t) - 1| > |\Re A(\pi/4) - 1| > e^{\sqrt{2}/(2a^{1/n})} \cos(\frac{\sqrt{2}}{2a^{1/n}}) - 1 \ge 1 - \cos(a^{-1/n}).$$

These inequalities and the symmetry imply the claimed lower bound on |A(t) - 1|. The second claim follows.

We note that when $a = 2^{o(n)}$ and n is sufficiently large, then the lower bound on $|f_{n,a}(z)|$ may be replaced by $|f_{n,a}(z)| > a/3$. To use Rouché's theorem, it is then enough that a, ρ and $||P||_1$ satisfy $a > (3e^{\rho}||P||_1)^{\frac{1}{1-\rho}}$. We can now derive upper bounds on the norms of $V_{g_{n,a}}$ and its inverse.

Lemma 3.8. For any $a > (||P||_1 \cdot C^{-1} \cdot e^{\rho})^{\frac{1}{1-\rho}}$ with $C = |1 - \cos(a^{-1/n}) - \frac{2e^{a^{-1/n}}}{na^{2/n}}|$, we have: $||V_{g_{n,a}}||_{\mathsf{F}} \le ane \quad and \quad ||V_{g_{n,a}}^{-1}||_{\mathsf{F}} \le n^{5/2}(||P||_1 + 1)a^{1/n}e^2.$

Proof. Let $\alpha_j = a^{1/n} e^{2i\pi j/n}$ be the roots of $f_{n,a}$ (for $0 \le j < n$). Thanks to the assumptions and Lemma 3.7, Theorem 3.4 allows us to locate the roots $(\beta_j)_{0\le j< n}$ of $g_{n,a}$ within distance 1/n from the α_j 's. Up to renumbering, we have $|\alpha_j - \beta_j| \le 1/n$ for all j. In particular, this implies that $|\beta_j| \le a^{1/n} + 1/n$ for all j. The first claim follows from (3.1).

Another consequence is that any power less than n of any $|\beta_j|$ is $\leq ae$. We start the estimation of $||V_{g_{n,a}}^{-1}||$ by considering the numerators in (3.2). Let $k_0 = 1 + \lfloor n(1-\rho) \rfloor$. For any $k < k_0$, we know that $E_k(\beta) = 0$. Using (3.3), we obtain $|e_k(\overline{\beta}^j)| = |\beta_j|^k \leq ae$ for $k < k_0$ and that $e_{k_0-1}(\overline{\beta}^j) =$ $(-1)^{k_0-1}\beta_j^{k_0-1}$. Then (3.3) gives $E_{k_0}(\beta) = (-1)^{k_0}p_{n-k_0} = (-1)^{k_0-1}\beta_j^{k_0} + e_{k_0}(\overline{\beta}^j)$, which implies that $|e_{k_0}(\overline{\beta}^j)| \leq |\beta_j|^{k_0} + |p_{n-k_0}|$. By induction, we obtain, for all $k < n - k_0$:

$$|e_{k_0+k}(\overline{\beta}^{j})| \leq |p_{n-k_0-k}| + |p_{n-k_0-k+1}\beta_j| + \dots + |p_{n-k_0}\beta_j^{k}| + |\beta_j|^{k_0+k}$$

$$\leq (||P||_1 + 1) \max(1, |\beta_j|^n),$$

so that $|e_k(\overline{\beta}^j)| \leq (||P||_1 + 1)ae$ for $k \geq 1$.

We now derive a lower bound on the denominators in (3.2). The separation of the β_j 's is close to that of the α_j 's. Concretely: $|\beta_j - \beta_k| \ge |\alpha_j - \alpha_k| - 2/n$ for all j, k. Therefore, we have $\prod_{k \ne j} |\beta_j - \beta_k| \ge \prod_{k \ne j} (|\alpha_j - \alpha_k| - 2/n)$. Using the identity $|\alpha_j - \alpha_k| = 2a^{1/n} \sin(|k-j|\pi/n)$ and elementary calculus, we obtain $\prod_{k \ne j} |\beta_j - \beta_k| \ge a^{\frac{n-1}{n}}/(ne)$. Indeed, recall that $\prod_{k \ne j} |\beta_j - \beta_k| \ge \prod_{k \ne j} (|\alpha_j - \alpha_k| - 2/n)$, and that $|\alpha_j - \alpha_k| = 2a^{1/n} \sin(|k-j|\pi/n)$. Standard bounds on the sine function give that $\sin(k\pi/n) \ge 2k/n$

for $1 \le k \le n/2$, and $\sin(k\pi/n) \ge 2 - 2k/n$ for $n/2 < k \le n$. We derive that:

$$\prod_{k \neq j} |\beta_j - \beta_k| \ge \prod_{k \neq j} |\alpha_j - \alpha_k| \cdot \prod_{\substack{k \neq j \\ |k-j| \le n/2}} \left(1 - \frac{1}{2a^{1/n}|k-j|} \right)^2$$
$$\ge |f'_{n,a}(\alpha_j)| \cdot \exp\left(2\sum_{1 \le k' \le n/2} \log\left(1 - \frac{1}{2a^{1/n}k'}\right) \right).$$

We have $\log(1 - \frac{1}{2a^{1/n}k'}) \ge \frac{-1}{a^{1/n}k'}$, and from the asymptotic expression of harmonic numbers, we can write $\sum_{k'=1}^{n/2} 1/k' \le \log(n/2) + 1$. We obtain:

$$\prod_{k \neq j} |\beta_j - \beta_k| \ge n a^{(n-1)/n} \cdot \left(\frac{ne}{2}\right)^{-2a^{-1/n}} \ge a^{(n-1)/n}/(ne)$$

Thus any coefficient w_{ij} of $V_{g_{n,a}}^{-1}$ satisfies $|w_{ij}| \le n(||P||_1 + 1)a^{1/n}e^2$. The claim follows from equivalence between the row and Frobenius norms.

We now assume that the p_j 's and a are integers. The following lemma states that, for a prime and sufficiently large, the polynomial $g_{n,a}$ is irreducible, and thus defines a number field.

Lemma 3.9. Assume that P is an integer polynomial. For any prime $a > ||P||_1 + 1$, the polynomial $g_{n,a}$ is irreducible over \mathbb{Q} .

Proof. Let β be a root of $g_{n,a}$. Then we have $a = |\beta^n + P(\beta)| \le |\beta|^n + ||P||_1 \max(1, |\beta|^n)$. The assumption on a implies that $|\beta| > 1$. In other words, all the roots of $g_{n,a}$ have a magnitude > 1. Now, assume by contradiction that $g_{n,a} = h_1h_2$ for some rational polynomials h_1, h_2 . Since $g_{n,a}$ is monic, it is primitive and we can choose h_1, h_2 as integer polynomials. The product of their constant coefficients is then the prime a. Hence the constant coefficient of h_1 or h_2 is ± 1 , which contradicts the fact that the roots of $g_{n,a}$ have magnitude > 1.

Overall, we have proved the following result.

Theorem 3.5. Let $\rho \in (0,1)$ and $p_j \in \mathbb{Z}$ for $1 \leq j \leq \rho \cdot n$. Then for $a \geq (3e^{\rho} ||P||_1)^{1/(1-\rho)}$ smaller than $2^{o(n)}$ and prime, and n sufficiently large, the polynomial $g_{n,a} = x^n + \sum_{1 \leq j \leq \rho \cdot n} p_j x^j + a$ is irreducible over \mathbb{Q} and satisfies:

$$\|V_{g_{n,a}}\|_{\mathsf{F}} \leq ane \quad and \quad \|V_{g_{n,a}}^{-1}\|_{\mathsf{F}} \leq n^{5/2} (\|P\|_1 + 1) a^{1/n} e^2.$$

In particular, if a and $||P||_1$ are polynomial in n, then both $||V_{g_{n,a}}||_{\mathsf{F}}$ and $||V_{g_{n,a}}^{-1}||_{\mathsf{F}}$ are polynomial in n.

3.7.4 Other "good" families of polynomials

We consider polynomials as $f_{n,\epsilon_0,\epsilon_1} = x^n + \varepsilon_1 \cdot x + \varepsilon_0$ for $\varepsilon_i \in \{\pm 1\}$. Notice that this class of polynomials includes the polynomials used in [BCLvV16]. Recall that $V_{f_{n,\epsilon_0,\epsilon_1}}$ denotes the Vandermonde matrix associated to $f_{n,\epsilon_0,\epsilon_1}$. We prove the following result.

Lemma 3.10. For every n > 2 and any $\varepsilon_0, \varepsilon_1 \in \{\pm 1\}$, we have:

$$\|V_{f_{n,\epsilon_0,\epsilon_1}}\|_{\mathsf{F}} \le 2n$$
 and $\|V_{f_{n,\epsilon_0,\epsilon_1}}^{-1}\|_{\mathsf{F}} \le 6n^{7/2}$.

We first use a general result on lacunary polynomials to estimate the magnitudes of the roots.

Proposition 3.1 ([Mig00, Thm. 1]). For any positive integer n and $1 \le k < n - 1$, let $P(x) = x^n + a_{n-k-1}x^{n-k-1} + \cdots + a_0$ be a complex polynomial, such that $a_0 \ne 0$. For any root α of P, we have

$$|\alpha| \le (n-k)^{\frac{1}{k+1}} \cdot \max_{1 \le j \le n} |a_{n-j}|^{1/j}.$$

Proof of Lemma 3.10. In our case, we see that any root α of $f_{n,\epsilon_0,\epsilon_1}$ is less than $2^{\frac{1}{n-1}}$. We use this observation several times below. Thanks to Equation (3.1), this gives that $\|V_{f_{n,\epsilon_0,\epsilon_1}}\|_{\mathsf{F}} \leq 2n$.

We use (3.2) to estimate $\|V_{f_{n,\epsilon_0,\epsilon_1}}^{-1}\|_{\mathsf{F}}$. From (3.3), we get that $|e_i(\overline{\alpha}^j)| = |\alpha_j|^i$ for $i \leq n-2$ and $j \leq n$, and $|e_{n-1}(\overline{\alpha}^j)| = |\varepsilon_0 - \alpha_j \cdot e_{n-2}(\overline{\alpha}^j)| \leq 3$. We now study the denominators of (3.2), that we can rewrite as $f'_{n,\epsilon_0,\epsilon_1}(\alpha_j) = \frac{\alpha_j(1-n)\varepsilon_1 - n\varepsilon_0}{\alpha_j}$. Using the triangle inequality, we have $|\alpha_j(1-n)\varepsilon_1 - n\varepsilon_0| \geq n - (n-1) \cdot 2^{\frac{1}{n-1}}$. Since the function $g(x) = (1+1/x)^x$ is strictly increasing, so is the sequence $a_n = (1 + \frac{n+1}{n^2})^{\frac{n}{n+1}}$. This gives that $a_n^{1-1/n^2} = (1 + \frac{n+1}{n^2})^{n-1} \geq 2$ for any $n \geq 3$. It follows that $n - (n-1) \cdot 2^{\frac{1}{n-1}} \geq 1/n^2$ for any $n \geq 3$. We conclude by observing that $|\alpha_j| < 2$ implies that $|f'(\alpha_j)| \geq \frac{1}{2n^2}$ and then $|w_{ij}| \leq 6n^2$. Equivalence between row and Frobenius norms gives the claim. \square

In this situation, $f_{n,\epsilon_0,\epsilon_1}$ may not be irreducible over \mathbb{Q} . For example, if $n \equiv 2 \mod 3$ and $\varepsilon_0 = \varepsilon_1 = 1$, then the primitive third roots of unity are also roots of $f_{n,1}$, hence $x^2 + x + 1$ is a factor. A similar situation occurs with $x^2 - x + 1$ if $n \equiv 2 \mod 6$ and $\varepsilon_0 = 1, \varepsilon_1 = -1$. This does not, however, impact the estimation of the norms.

3.8 On small elements and $f'(\alpha)$

In Section 3.7.1, we discussed the possibility to use $f'(\alpha)$ for reductions between dual (resp. primal) RLWE and primal-RLWE (resp. PLWE), as it is the case that $f'(\alpha) \in \mathcal{C}_{\mathcal{O}} \cap (\mathcal{O}_{K}^{\vee})^{-1}$. The results of Section 3.6 are meaningful for our applications when there are smaller elements in $(\mathcal{O}_{K}^{\vee})^{-1}$ and $\mathcal{C}_{\mathcal{O}}$ than in the ideal generated by $f'(\alpha)$. More formally, we show that there are fields K for which

$$\lambda_1((\mathcal{O}_K^{\vee})^{-1}) < \lambda_1((f'(\alpha))) \quad (\text{ resp. } \lambda_1(\mathcal{C}_{\mathcal{O}}) < \lambda_1((f'(\alpha)))).$$

By Lemma 2.13, it suffices that $\Delta_f > \Delta_K^{3/2}$ (resp. $\Delta_f > \mathcal{N}_{\mathcal{O}_K}(\mathcal{C}_{\mathcal{O}})\Delta_K^{1/2}$). Below, we give a family of number fields K of degree 3 with defining polynomials f such that $f'(\alpha)$ can have an arbitrarily large algebraic norm, relatively to those of $(\mathcal{O}_K^{\vee})^{-1}$ and $\mathcal{C}_{\mathcal{O}}$.

Lemma 3.11. Let $q \neq 3$ be a prime integer such that $q^2 \not\equiv 1 \mod 9$. Let $f = x^3 - q^2$, $K = \mathbb{Q}[x]/f$ and $\mathcal{O} = \mathbb{Z}[x]/f \simeq \mathbb{Z}[\alpha]$.

- 1. We have $\mathcal{N}(f'(\alpha)) = \Delta_f = 3^3 \cdot q^4$ and $\mathcal{N}((\mathcal{O}_K^{\vee})^{-1}) = \Delta_K = 3^3 \cdot q^2$.
- 2. If $\mathcal{C}_{\mathcal{O}}$ is the conductor of \mathcal{O} , then $\mathcal{N}_{\mathcal{O}}(\mathcal{C}_{\mathcal{O}}) = [\mathcal{O}_K : \mathcal{O}] = q$ and $\mathcal{N}_{\mathcal{O}_K}(\mathcal{C}_{\mathcal{O}}) = q^2$.

The family of f's considered in Lemma 3.11 is restrictive. Numerical experiments suggest that polynomials $f = x^p - q^2$ with p, q distinct primes and $q^2 \not\equiv 1 \mod p^2$ give $[\mathcal{O}_K : \mathcal{O}] = \mathcal{N}_{\mathcal{O}}(\mathcal{C}_{\mathcal{O}}) = q^{\frac{p-1}{2}}$ and $\mathcal{N}_{\mathcal{O}_K}(\mathcal{C}_{\mathcal{O}}) = q^{p-1}$.

Proof. A determinant computation gives $\Delta_f = \operatorname{Res}(f, f') = 3^3 \cdot q^4$. From this factorization and the formula $\Delta_f = [\mathcal{O}_K : \mathcal{O}]^2 \cdot \Delta_K$, we can deduce that 3 and q are the only possible prime factors of $[\mathcal{O}_K : \mathcal{O}]$. It is known (see, e.g., [Ste17, p.48]) that a prime integer p divides this index if and only if there is at least one prime \mathcal{O} -ideal factor of $p\mathcal{O}$ which is not invertible as an \mathcal{O} -ideal. This property amounts to checking divisibility between polynomials (Kummer-Dedekind's theorem, [Ste17, Thm. 3.1, p.31]), and \mathcal{O} is said to be singular over p.

We first show that \mathcal{O} is not singular over 3 but is singular over q. The reduction of f modulo 3 is $x^3 - 1 = (x - 1)^3$ in \mathbb{F}_3 . Division of f by x - 1 gives $f = (x - 1)(x^2 + x + 1) + 1 - q^2$, so from the assumptions on q, 3² does not divide the remainder $1 - q^2$. This precisely means that \mathcal{O} is not singular over 3, and we deduce that 3 divides Δ_K . On the other hand, the reduction of f modulo q is x^3 in \mathbb{F}_q . Division of f by x gives $f = x \cdot x^2 - q^2$, so that q^2 divides the remainder: the order \mathcal{O} is singular over q. In particular, the index $[\mathcal{O}_K : \mathcal{O}]$ is either q or q^2 .

From the factorization of f modulo q, we also know that the ideal $\mathfrak{p}_q = \langle q, \alpha \rangle$ is the only prime in \mathcal{O} containing $q\mathcal{O}$, and that it is not invertible. From [Ste17, ex. 25, p. 53], this also means that $\mathcal{C}_{\mathcal{O}} \subseteq \mathfrak{p}_q$, where $\mathcal{C}_{\mathcal{O}}$ is the (non-trivial) conductor of \mathcal{O} .

Using [Ste17, Cor. 3.2, p. 32], we know that $\beta := \frac{1}{q}\alpha^2$ is not in \mathcal{O} . One checks that the minimal polynomial of β over \mathbb{Q} is $x^3 - q$, hence $\beta \in \mathcal{O}_K$. In particular, we have a ring extension $\mathcal{O} \subseteq \mathcal{O}[\beta] \subseteq \mathcal{O}_K$. Observe that $\mathbb{Z}[\beta]$ is regular above q: reducing $x^3 - q$ modulo q gives again x^3 , but the remainder by division by x is now q. Now, the order $\mathcal{O}[\beta]$ is a common extension of \mathcal{O} and $\mathbb{Z}[\beta]$, and from [Ste17, Le. 3.8, p. 33], ring extensions do not add new singular primes. This implies that $\mathcal{O}[\beta]$ is a Dedekind ring in \mathcal{O}_K . Moreover, from [Ste17, Le. 3.20, p. 39], we get that $\mathcal{O}[\beta] = \mathcal{O}_K$. We also obtain that $q\mathcal{O}_K \subseteq \mathfrak{P}_q := \langle q, \beta \rangle = \beta \mathcal{O}_K$.

We first observe that $\beta^2 - \alpha = 0$, which means that $\mathcal{O}[\beta] = \{\lambda\beta + \mu : \lambda, \mu \in \mathcal{O}\}$. We readily check that $q(\lambda\beta + \mu)$ and $\alpha(\lambda\beta + \mu)$ are elements in \mathcal{O} for any $\lambda, \mu \in \mathbb{Z}[\alpha]$, so we actually have that $\mathfrak{p}_q := \langle q, \alpha \rangle \mathcal{O} = \mathcal{C}_{\mathcal{O}}$. This means that $\mathcal{O}/\mathfrak{p}_q \simeq \mathbb{F}_q$ or, equivalently, that $\mathcal{N}_{\mathcal{O}}(\mathcal{C}_{\mathcal{O}}) = q$. We now show that $|\mathcal{O}[\beta]/\mathcal{O}| = [\mathcal{O}:\mathcal{C}_{\mathcal{O}}]$, where the left cardinality is taken for the quotient of the additive groups. Now two elements $\lambda\beta + \mu, \lambda'\beta + \mu'$ are in the same class if and only if $(\lambda - \lambda')\beta$ is in \mathcal{O} . This amounts to asking that $\lambda - \lambda' \in \mathcal{C}_{\mathcal{O}}$, so that the classes of the quotient ring $\mathcal{O}/\mathcal{C}_{\mathcal{O}}$ are in one-to-one correspondence with the classes of the quotient group $\mathcal{O}[\beta]/\mathcal{O}$. In other words, we have $[\mathcal{O}_K:\mathcal{O}] = q$.

We now describe $\mathcal{C}_{\mathcal{O}}$ as an \mathcal{O}_K -ideal. Since $\beta^2 = \alpha$, we have $\mathcal{C}_{\mathcal{O}} \subsetneq \mathfrak{P}_q = \beta \mathcal{O}_K$ as \mathcal{O}_K -ideals. On the other hand, we have $\mathfrak{P}_q^2 = \beta^2 \mathcal{O}_K = \alpha \mathcal{O}_K \subseteq \mathcal{C}_{\mathcal{O}}$ as \mathcal{O}_K -ideals. As \mathfrak{P}_q is prime in \mathcal{O}_K , we get $\mathcal{C}_{\mathcal{O}} = \mathfrak{P}_q^2$. We now obtain that $\mathcal{N}_{\mathcal{O}_K}(\mathcal{C}_{\mathcal{O}}) = \mathcal{N}_{\mathcal{O}_K}(\mathfrak{P}_q^2) = q^2$.

3.9 Search to decision dual-RLWE

The reduction relies on the recent technique of [PRSD17]. To leverage it, we use a generalized Leftover Hash Lemma over rings. The proof generalizes a technique used in [SS11] to the case where the irreducible factors of the defining polynomial (of K) reduced modulo q do not share the same degree. Alternatively, a generalization of the regularity lemma from [LPR13, Se. 7] to arbitrary number fields could be used. Such a generalization may go through and improve our results a little.

3.9.1 A ring-based Leftover Hash Lemma

Let $m \geq 2$. We identify any rank $m \mathcal{O}_K$ -module $M \subseteq K^m$ with the lattice $\sigma(M) \subseteq H^m$. For such modules, the dual may be defined as

$$\widehat{M} = \{ \mathbf{t} \in K^m : \forall \mathbf{x} \in M, \operatorname{Tr}(\langle \mathbf{t}, \mathbf{x} \rangle) \in \mathbb{Z} \}.$$

Here $\langle \cdot, \cdot \rangle$ is the *K*-bilinear map defined by $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^{m} x_i y_i$. We have $\sigma(\widehat{M}) = \overline{\sigma(M)^*}$ in H^m . For some $q \geq 2$ and a fixed $\mathbf{a} \in (\mathcal{O}_K/q\mathcal{O}_K)^m$, we focus on the modules:

$$L(\mathbf{a}) = \frac{\mathbf{a}}{q} \mathcal{O}_K^{\vee} + (\mathcal{O}_K^{\vee})^m \text{ and } \mathbf{a}^{\perp} = \{ \mathbf{t} \in \mathcal{O}_K^m : \langle \mathbf{t}, \mathbf{a} \rangle = 0 \mod q \mathcal{O}_K \}.$$

To prove our Leftover Hash Lemma variant, the main argument relies on an estimation of $\lambda_1^{\infty}(\widehat{\mathbf{a}}^{\perp})$, which is obtained by combining the following two lemmas. The first one was stated in [LS15, Se. 5]

without a proof, for the case of cyclotomic fields (this restriction is unnecessary). For the sake of completeness, we give its proof here.

Lemma 3.12. Let $q \ge 2$ and $\mathbf{a} \in (\mathcal{O}_K/q\mathcal{O}_K)^m$. Then we have $\widehat{\mathbf{a}^{\perp}} = L(\mathbf{a})$.

Proof. We proceed by double inclusion, starting with $L(\mathbf{a}) \subseteq \widehat{\mathbf{a}^{\perp}}$. Let $\mathbf{x} = (x_1, \ldots, x_m) \in \mathbf{a}^{\perp}$ and $\mathbf{t} = (t_1, \ldots, t_m) \in L(\mathbf{a})$. By definition, there exist $s \in \mathcal{O}_K^{\vee}$ and $b_1, \ldots, b_m \in \mathcal{O}_K^{\vee}$ such that $t_i = \frac{a_i}{q}s + b_i$, for all *i*. Then the element $\operatorname{Tr}(\langle \mathbf{t}, \mathbf{x} \rangle) = \frac{1}{q} \operatorname{Tr}(s\langle \mathbf{a}, \mathbf{x} \rangle) + \sum_{i=1}^{m} \operatorname{Tr}(x_i b_i)$ is an integer. Indeed, by definition of \mathbf{x} , the product $s\langle \mathbf{a}, \mathbf{x} \rangle$ belongs to $q\mathcal{O}_K^{\vee}$. This implies that all traces are rational integers, which completes the proof of the first inclusion.

By duality, the reverse inclusion is equivalent to $\widehat{L}(\mathbf{a}) \subseteq \mathbf{a}^{\perp}$. Let $\mathbf{y} \in \widehat{L}(\mathbf{a})$. As $\frac{\mathbf{a}}{q} \in L(\mathbf{a})$ we obtain that $\operatorname{Tr}(\langle \mathbf{y}, \mathbf{a} \rangle) \in q\mathbb{Z}$. This implies that we have $\operatorname{Tr}(\langle \mathbf{y}, \mathbf{b} \rangle) \in \mathbb{Z}$ for all $\mathbf{b} \in (\mathcal{O}_K^{\vee})^m$. Taking for \mathbf{b} vectors with one coordinate arbitrary in \mathcal{O}_K^{\vee} and 0 for the rest, we see that all y_i 's belong to $\mathcal{O}_K^{\vee\vee} = \mathcal{O}_K$, hence $\mathbf{y} \in \mathcal{O}_K^m$. The fact that $\operatorname{Tr}(\langle \mathbf{y}, \mathbf{b} \rangle) \in \mathbb{Z}$ for all $\mathbf{b} \in (\mathcal{O}_K^{\vee})^m$ also implies that $\operatorname{Tr}(s\langle \frac{\mathbf{a}}{q}, \mathbf{y} \rangle)$ is an integer for all $s \in \mathcal{O}_K^{\vee}$, so that $\langle \frac{\mathbf{a}}{q}, \mathbf{y} \rangle \in \mathcal{O}_K^{\vee\vee} = \mathcal{O}_K$. Equivalently, we have $\mathbf{y} \in \mathbf{a}^{\perp}$.

We now obtain a probabilistic lower bound on $\lambda_1^{\infty}(\widehat{\mathbf{a}^{\perp}}) = \lambda_1^{\infty}(L(\mathbf{a}))$. In full generality, it should depend on the ramification of the selected prime integer q, *i.e.*, the exponents appearing in the factorization of $q\mathcal{O}_K$ in prime ideals. It is a classical fact that the ramified prime integers are exactly the primes dividing the discriminant of the field, so that there are only finitely many such q's. Moreover, it is always possible to use modulus switching techniques ([BLP+13, LS15]) if q ramifies. Therefore, we consider only the non-ramified case.

Lemma 3.13. Let $q \ge 2$ a prime that does not divide Δ_K . For any $m \ge 2$ and $\delta > 0$, and except with a probability $\le 2^{3n(m+1)}q^{-mn\delta}$ over the uniform choice of $\mathbf{a} \in ((\mathcal{O}_K/q\mathcal{O}_K)^{\times})^m$, we have:

$$\lambda_1^{\infty}(L(\mathbf{a})) \ge \Delta_K^{-1/n} \cdot q^{-\frac{1}{m}-\delta}$$

Proof. Thanks to the assumption on q, we can write $q\mathcal{O}_K = \mathfrak{p}_1 \dots \mathfrak{p}_k$ for distinct prime ideals \mathfrak{p}_i . By Lemma 3.6 and the Chinese Remainder Theorem, we have $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee} \simeq \mathcal{O}_K/q\mathcal{O}_K \simeq \bigoplus_{i=1}^k \mathbb{F}_{q^{d_i}}$, where $q^{d_i} = \mathcal{N}(\mathfrak{p}_i)$.

Let $\mathbf{a} = (a_1, \ldots, a_m)$ sampled uniformly in $((\mathcal{O}_K/q\mathcal{O}_K)^{\times})^m$. Fix some bound B > 0 and let p_B be the probability that $qL(\mathbf{a}) = \mathbf{a}\mathcal{O}_K^{\vee} + q(\mathcal{O}_K^{\vee})^m$ contains a $\mathbf{t} = (t_1, \ldots, t_m)$ such that $0 < \|\mathbf{t}\|_{\infty} < B$. Our goal is to bound p_B from above. By the union bound, we have that

$$p_B \leq \sum_{s \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}} \sum_{\substack{\mathbf{t} \in (\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee})^m \\ 0 < \|\mathbf{t}\|_{\infty} < B}} p(\mathbf{t}, s),$$

with $p(\mathbf{t}, s) = \Pr_{\mathbf{a}}[\forall j, t_j = a_j s \mod q\mathcal{O}_K^{\vee}]$ for any s and \mathbf{t} over $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$. By independance of the a_j 's, we can write $p(\mathbf{t}, s) = \prod_{j \in [m]} p(t_j, s)$ with $p(t_j, s) = \Pr_{a_j}[t_j = a_j s \mod q\mathcal{O}_K^{\vee}]$. As $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ and $\mathcal{O}_K/q\mathcal{O}_K$ are isomorphic, estimating this probability amounts to studying the solutions in $(\mathcal{O}_K/q\mathcal{O}_K)^{\times}$ of the equation $t = as \mod q\mathcal{O}_K$, for all $t, s \in \mathcal{O}_K/q\mathcal{O}_K$.

Note that if there is an i such that $t = 0 \mod \mathfrak{p}_i$ and $s \neq 0 \mod \mathfrak{p}_i$, or vice-versa, then there is no solution, so that p(t, s) = 0. Now, assume that s and t are 0 modulo the same \mathfrak{p}_i 's. Let $S \subseteq [k]$ denote the set of their indices, and let d_S be such that $q^{d_S} = \mathcal{N}(\prod_{i \in S} \mathfrak{p}_i)$. On the one hand, for all $i \in [k] \setminus S$, both t and s are invertible modulo \mathfrak{p}_i so there is exactly one solution modulo those i's. On the other hand, for all $i \in S$, all the elements of $\mathbb{F}_{q^{d_i}}^{\times}$ are solutions. This gives $\prod_{i \in S} (q^{d_i} - 1)$ possibilities out of the $\prod_i (q^{d_i} - 1)$ elements of $(\mathcal{O}_K/q\mathcal{O}_K)^{\times}$. Overall, we obtain that $p(t, s) = \prod_{i \in [k] \setminus S} (q^{d_i} - 1)^{-1}$. Hence,

for **t** with coordinates t_j such that s and all t_j 's are 0 modulo the same \mathbf{p}_i 's, we have:

$$p(\mathbf{t},s) = q^{-m(n-d_S)} \prod_{i \in [k] \setminus S} (1 - \frac{1}{q^{d_i}})^{-m} \le q^{-m(n-d_S)} \cdot 2^{mk},$$

the last inequality coming from the fact that $1 - 1/q^{d_i} \ge 1/2$ for all *i*.

Let τ denote the isomorphism mapping $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ to $\mathcal{O}_K/q\mathcal{O}_K$. The probability to bound is now

$$p_B \leq 2^{mk} \cdot \sum_{S \subseteq [k]} \sum_{\substack{\tau(s) \in \mathcal{O}_K/q\mathcal{O}_K \\ \forall i \in S: \mathfrak{p}_i \mid \tau(s)}} \sum_{\substack{\tau(t) \in (\mathcal{O}_K/q\mathcal{O}_K)^m \\ 0 < \|\mathbf{t}\|_{\infty} < B \\ \forall j, \forall i \in S: \mathfrak{p}_i \mid \tau(t_j)}} q^{-m(n-d_S)}.$$

For any r > 0, we let $\mathcal{B}(r)$ denote the (open) ball in H of center 0 and radius r, with respect to the infinity norm. Such a ball has a volume $\operatorname{Vol}(\mathcal{B}(r)) = (2r)^n$. For any $S \subseteq [k]$, we define $N(B,S) = |\mathcal{B}(B) \cap \mathcal{L}(\tau^{-1}(\prod_{i \in S} \mathfrak{p}_i))| - 1$. Since there are 2^k subsets in [k] and q^{n-d_S} elements $\tau(s) \in \mathcal{O}_K/q\mathcal{O}_K$ such that $\mathfrak{p}_i|s$ for all $i \in S$, we have

$$p_B \le 2^{k(m+1)} \cdot \max_{S \subseteq [k]} \frac{N(B,S)^m}{q^{(n-d_S)(m-1)}}.$$
(3.4)

We now give an upper bound for N(B, S), from which we will obtain the result. Let $I_S = \prod_{i \in S} \mathfrak{p}_i$ and $\lambda_S = \lambda_1^{\infty}(\tau^{-1}(I_S))$. Observe that any two distinct balls of radius $\lambda_S/2$ and centered around elements of $\mathcal{B}(B) \cap \mathcal{L}(\tau^{-1}(I_S))$ do not intersect. Moreover, all of them are contained in $\mathcal{B}(B + \lambda_S/2)$. This implies that

$$N(B,S) \le \frac{\operatorname{Vol}(\mathcal{B}(B+\lambda_S/2))}{\operatorname{Vol}(\mathcal{B}(\lambda_S/2))} = \left(\frac{2B}{\lambda_S} + 1\right)^n.$$

It remains to give a lower bound on λ_S . As $\tau^{-1}(I_S) = I_S \mathcal{O}_K^{\vee}$, we have $\mathcal{N}(\tau^{-1}(I_S)) = q^{d_S}/\Delta_K$. With Lemma 2.13, this gives $\Delta_K^{-1/n} q^{d_S/n} \leq \lambda_S$. If we set $B = \Delta_K^{-1/n} q^{\beta}$, then $n\beta < d_S$ leads to N(B, S) = 0 and $n\beta \geq d_S$ implies the upper bound $N(B, S) \leq 2^{2n} q^{n\beta-d_S}$. With (3.4), this gives

$$p_B \le 2^{(m+1)(k+2n)} \cdot \max_{\substack{S \subseteq [k] \\ d_S \le n\beta}} q^{m(\beta-1)n+(n-d_S)}.$$

The maximum is reached for $d_S = 0$ (*i.e.*, when $S = \emptyset$). In this case, the exponent of q is $-mn\delta$ for $\beta = 1 - \frac{1}{m} - \delta$. We obtain that $\lambda_1^{\infty}(qL(\mathbf{a})) \ge \Delta_K^{-1/n} q^{1-\frac{1}{m}-\delta}$ except with probability $\le 2^{3n(m+1)}q^{-mn\delta}$. \Box

We are now ready to state the variant of the Leftover Hash Lemma.

Theorem 3.6. Let $q \ge 2$ prime that does not divide Δ_K . Let $\delta > 0, \epsilon \in (0, 1/2)$ and $m \ge 2$. For a given **a** in $((\mathcal{O}_K/q\mathcal{O}_K)^{\times})^m$, let U_a be the distribution of $\sum_{i\le m} t_i a_i$ where the vector $\mathbf{t} = (t_1, \ldots, t_m)$ is sampled from $D_{\mathcal{O}_K,s}$ with $s \ge \sqrt{\log(2mn(1+1/\epsilon))/\pi} \cdot \Delta_K^{1/n} q^{1/m+\delta}$. Then, except for $\le 2^{3n(m+1)}q^{-mn\delta}$ of **a**'s, the distance to uniformity of $U_{\mathbf{a}}$ is $\le 2\epsilon$.

Proof. First we note that the map $\mathbf{t} \mapsto \sum_{i \leq m} t_i a_i$ is a well-defined surjective \mathcal{O}_K -module homomorphism from \mathcal{O}_K^m to $\mathcal{O}_K/q\mathcal{O}_K$, with kernel \mathbf{a}^{\perp} . The distance to uniformity of $U_{\mathbf{a}}$ is hence the same as the distance to uniformity of $\mathbf{t} \mod \mathbf{a}^{\perp}$. By Lemma 2.11, the claim follows whenever $s \geq \eta_{\epsilon}(\mathbf{a}^{\perp})$. By Lemma 2.9, it suffices to find an appropriate lower bound on $\lambda_1^{\infty}(L(\mathbf{a}))$. Lemma 3.13 allows to complete the proof.

Corollary 3.2 (Leftover Hash lemma). If t is sampled from $D_{\mathcal{O}_{K},s}$ with $s \geq \sqrt{\log(2mn(1+1/\epsilon))/\pi}$.

 $\Delta_K^{1/n} q^{1/m+\delta}$, and the a_i 's are sampled from $U((\mathcal{O}_K/q\mathcal{O}_K)^{\times})$, then:

$$\Delta \qquad \left[\left(a_1, \dots, a_m, \sum_{i \le m} t_i a_i \right), U \left(\left((\mathcal{O}_K/q\mathcal{O}_K)^{\times} \right)^m \times \mathcal{O}_K/q\mathcal{O}_K \right) \right] \\ \le 2\epsilon + 2^{3n(m+1)} \cdot q^{-mn\delta}.$$

3.9.2 Search RLWE to decision RLWE

We now give the reduction from search to decision. As all proofs can be done similarly, we focus on the dual-RLWE version of the problems. For the sake of simplicity, we consider only the case of diagonal covariance matrices. The proof readily extends to general covariance matrices. To obtain the reduction, we need to generate suitable new samples from a starting set of samples from search dual-RLWE. The lemma below is adapted from [LS15, Le. 4.15]. We will use it to analyze the error distribution we get when generating new samples.

Lemma 3.14. Let $\alpha > 0$, \mathcal{L} a rank- $\mathcal{m} \mathcal{O}_K$ -module, $\epsilon \in (0, 1/2)$, a vector $\mathbf{t} \in D_{\mathcal{L}+\mathbf{c},\mathbf{r}}$ for some $\mathbf{c} \in H^m$, and $e' \in K_{\mathbb{R}}$ chosen according to D_{α}^H . If $r_i \ge \eta_{\epsilon}(\mathcal{L})$ and $\frac{\alpha}{\delta_i} \ge \eta_{\epsilon}(\mathcal{L})$ for all i, then $\Delta(\langle \mathbf{t}, \mathbf{e} \rangle + e', D_{\mathbf{x}}^H) \le 4\epsilon$ with $x_i = \sqrt{(r_i\delta_i)^2 + \alpha^2}$ and $\delta_i = (\sum_{k \in [m]} |\sigma_i(e_k)|^2)^{1/2}$ for all i.

We can now give a reduction from search dual-RLWE to worst-case decision dual-RLWE. It may be combined with the worst-case decision dual-RLWE to decision dual-RLWE from Lemma 3.5.

Theorem 3.7. Let $\mathbf{r} \in (\mathbb{R}^{\geq 0})^n$ be such that $r_i = r_{i+s_2}$ for any $i > s_1$ and $r_i \leq r$ for some r > 0. Let $d = \sqrt{n} \cdot \Delta_K^{1/n} q^{1/m+1/n}$, and consider $\Sigma = \{\mathbf{r}' : r'_i \leq \sqrt{d^2 \cdot r^2 \cdot m + d^2}\}$. Then there exists a probabilistic polynomial-time reduction from search dual-RLWE_{q,Dr} with $m \leq q/(2n)$ input samples to worst-case decision dual-RLWE_{q,\Sigma}.

Proof. We have m samples $(a_i, b_i = a_i s + e_i) \in \mathcal{O}_K/q\mathcal{O}_K \times K_{\mathbb{R}}/q\mathcal{O}_K^{\vee}$ from the dual-RLWE distribution $\mathsf{A}_{q,\mathbf{r}}^{\vee}(s)$, for a uniform $s \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$ that we want to find. This is equivalent to finding the error term $\mathbf{e} = (e_1, \ldots, e_m)$. By assumption on m, the a_i 's are all invertible with non-negligible probability. If it is not the case, the reduction aborts. From now on, we hence assume that they are uniformly distributed in $(\mathcal{O}_K/q\mathcal{O}_K)^{\times}$.

We use the same technique as in [PRSD17], in that we find the *i*th embeddings $\sigma_i(e_1), \ldots, \sigma_i(e_m)$ of the error terms by constructing an *m*-dimensional instance of the Oracle Hidden Center Problem (OHCP). The only difference consists in the way we create the samples that we give to the decision oracle. The reduction uses the dual-RLWE decision oracle to build the oracles $\mathcal{O}_i : \mathbb{R}^m \times \mathbb{R}^{\geq 0} \to \{0, 1\}$ for $i \leq s_1$ and $\mathcal{O}_i : \mathbb{C}^m \times \mathbb{R}^{\geq 0} \to \{0, 1\}$ for $s_1 < i \leq s_1 + s_2$.

For $i \leq s_1$, we define $k_i : \mathbb{R} \to K_{\mathbb{R}}$ as $k_i(x) = \sigma^{-1}(x \cdot \mathbf{v}_i)$ and for $s_1 < i \leq s_1 + s_2$, we define $k_i : \mathbb{C} \to K_{\mathbb{R}}$ as $k_i(x) = \sigma^{-1}(x \cdot \mathbf{v}_i + \overline{x} \cdot \mathbf{v}_{i+s_2})$, where the \mathbf{v}_i 's form the canonical basis of H.

On input $(z_1, \ldots, z_m, \alpha)$, oracle \mathcal{O}_i will output 1 with probability depending on $\exp(\alpha) \|\mathbf{e} - \overline{\mathbf{z}}\|$, where $\overline{\mathbf{z}} = (k_i(z_1), \ldots, k_i(z_m))$. It works as follows. It first chooses a uniform $s' \in \mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$. On input $(z_1, \ldots, z_m, \alpha)$, it samples $\mathbf{t} = (t_1, \ldots, t_m) \in \mathcal{O}_K^m$ Gaussian with parameter $\exp(\alpha) \cdot \sqrt{n} \cdot \Delta_K^{1/n} q^{1/m+1/n}$ and some e' from D_d . The oracle then creates $(a', b') = (\langle \mathbf{t}, \mathbf{a} \rangle, \langle \mathbf{t}, \mathbf{b} - \overline{\mathbf{z}} \rangle + a's' + e')$, where $\mathbf{b} = (b_1, \ldots, b_m)$.

By Corollary 3.2, the distribution of $(\mathbf{a}, \langle \mathbf{t}, \mathbf{a} \rangle)$ is exponentially close to $U(((\mathcal{O}_K/q\mathcal{O}_K)^{\times})^m \times \mathcal{O}_K/q\mathcal{O}_K)$. Since $b_j = a_j s + e_j$ for all j, we get $b' = a'(s + s') + \langle \mathbf{t}, \mathbf{e} - \bar{\mathbf{z}} \rangle + e'$, so oracle \mathcal{O}_i creates RLWE samples for a uniformly distributed s + s', provided the error term follows a suitable distribution. We let $\delta_\ell = (\sum_{j \in [m]} \sigma_\ell(e_j - k_i(z_j))|^2)^{1/2}$ for $\ell \leq n$. In particular, we have $\delta_i = \|\sigma_i(e_1) - z_1, \ldots, \sigma_i(e_m) - z_m\|$. Let us now study the distribution of the error term $\langle \mathbf{t}, \mathbf{e} - \bar{\mathbf{z}} \rangle + e'$. We can see that once the value of $\langle \mathbf{t}, \mathbf{a} \rangle = c$ and the a_i 's are known, one can write $\mathbf{t} = (ca_1^{-1}, 0, \ldots, 0) + (-a_1^{-1}\sum_{i>2} t_i a_i, t_2, \ldots, t_m)$, where the second vector belongs to \mathbf{a}^{\perp} . This means that the actual support of \mathbf{t} is a shift of the \mathbf{a}^{\perp} lattice by the vector $(ca_1^{-1}, 0, \dots, 0)$. Using Lemma 3.14, we get that the distribution of the error is $D_{\mathbf{x}}^H$ where $x_i = \sqrt{\exp^2(\alpha) \cdot d^2 \cdot \delta_i^2 + d^2}$.

Let $\mathcal{S}_{i,(z_1,\ldots,z_m,\alpha)}$ be the samples obtained by applying the procedure above many times. Oracle \mathcal{O}_i calls the dual-RLWE decision oracle with these and outputs 1 if and only if the latter accepts. With non-negligible probability over the choice of the initial errors, the distribution of the samples we get when we call the oracle \mathcal{O}_i on (0,0) belongs to the set Σ . One can now show that using the same technique as in [PRSD17], it is possible to recover good approximations of the vector $(\sigma_i(e_1),\ldots,\sigma_i(e_m))$. By substracting them from the initial search samples, rounding and then taking the inverses of the a_i 's, we obtain s.

3.10 On Vandermonde matrices and the expansion factor

In the study of algebraic variants of LWE, the expansion factor is an important parameter. For example, the PLWE to MP-LWE reduction from Chapter 4 requires that the expansion factor of the polynomial parameterizing PLWE be small. The polynomials f for which we managed to bound $||V_f||_{\mathsf{F}}$ and $||V_f^{-1}||_{\mathsf{F}}$ in Theorem 3.5 have small expansion factors.

In this section, we study the relationship between $\|V_f\|_{\mathsf{F}}, \|V_f^{-1}\|_{\mathsf{F}}$ and the expansion factor $\mathsf{EF}(f)$ of an arbitrary polynomial f. We first show that there exist polynomials f with small expansion factors but large $\|V_f^{-1}\|_{\mathsf{F}}$. Then, we show that if $\|V_f\|_{\mathsf{F}}$ and $\|V_f^{-1}\|_{\mathsf{F}}$ are both small, the expansion factor $\mathsf{EF}(f)$ will also be small.

For integers $n \ge 4, 2 \le k < n/2, a \ge 2$, consider the family of polynomials given by

$$g_{n,k,a} = x^n - 2(ax - 1)^k.$$

The factor 2 is used to ensure irreducibility by way of Eisenstein's criterion. Such polynomials have a "gap" in their coefficients. Considering a, k as function of n, their expansion factors are polynomially bounded if for example $a \leq poly(n)$ and k is constant, or if a is constant and $k \leq O(\log n)$.

Besides, Bugeaud and Mignotte showed that there is a cluster of k roots exponentially close to the real 1/a. In particular, if the other roots are not too far away from this cluster, the denominators in (3.2) force $\|V_f^{-1}\|_{\mathsf{F}}$ to be exponentially large. We adapt some results of [BM10]; in particular, we locate the roots outside the cluster to be at distance at most a from the origin. This enables us to prove that $\|V_f^{-1}\|_{\mathsf{F}}$ is exponentially large in n.

Lemma 3.15 (Adapted from [BM10]). If $(1 + 2^{1-n/k})^{n/k} < a$, then the polynomial $g_{n,k,a}$ has k roots in the disk $D(\frac{1}{a}, \frac{1}{a^{n/k}})$.

Proof. We apply Rouché's theorem. Write $g_{n,k,a} = f + P$, where $f = -2(ax - 1)^k$, and $P = x^n$ is the "perturbation." For any $z = \frac{1}{a} + \frac{e^{it}}{a^{n/k}}$ on the circle, we have $|f(z)| = \frac{2}{a^{n-k}}$ and $|P(z)| \le \left(\frac{1}{a} + \frac{1}{a^{n/k}}\right)^n$, so that the assumption gives |P(z)| < |f(z)|. We conclude using Theorem 3.4 and the fact that f has a root of multiplicity k in the disk.

Lemma 3.16. If $a > 4^{\frac{n+2k}{n-2k}}$, then the polynomial $g_{n,k,a}$ has all its roots in the disk $D(\frac{1}{a}, a^{\frac{n}{2(n-k)}} - \frac{1}{a^{n/k}})$.

Proof. Write $P = -2(ax-1)^k$ and $f = x^n$. For any z on the boundary of the disk, we have $|f(z)| \ge (a^{\frac{n}{2(n-k)}} - \frac{1}{a} - \frac{1}{a^{n/k}})^n \ge a^{\frac{n^2}{2(n-k)}} \cdot 2^{-n}$. If we write $P = \sum_i p_i x^i$, then $|p_i| = 2a^i {k \choose i}$ so that $||P||_1 = 2(a+1)^k$. We obtain

$$|P(z)| \le \max(1, |z|^k) \cdot ||P||_1 \le 2(a+1)^k \left(a^{\frac{n}{2(n-k)}} - \frac{1}{a} - \frac{1}{a^{n/k}}\right)^k$$

and the assumption implies that |P(z)| < |f(z)| on the boundary of the disk. We conclude using Rouché's theorem (Theorem 3.4).

The term " $-\frac{1}{a^{n/k}}$ " in the radius cancels in the next proof. As a consequence of these lemmata, we can show that the inverse Vandermonde associated to $g_{n,k,a}$ has several exponentially large entries.

Proposition 3.2. Let $n \ge 4, 2 \le k < n/2, a \ge 2$ be integers such that $a > \max\left((1+2^{1-n/k})^{n/k}, 4^{\frac{n+2k}{n-2k}}\right)$. Then $\|V_{g_{n,k,a}}^{-1}\|_{\infty} \ge \frac{a^{n/2-n/k}}{2^{k-1}}$.

Proof. The assumption on a allows us to apply the two lemmata above. Let $\alpha_1, \ldots, \alpha_k$ be the roots in the disk $D(\frac{1}{a}, \frac{1}{a^{n/k}})$ (their cardinality is provided by Lemma 3.15). We have, for all $i \leq k$, that $\prod_{j=1, j\neq i}^k |\alpha_i - \alpha_j| \leq \frac{2^{k-1}}{a^{n-n/k}}$. Let $\alpha_{k+1}, \ldots, \alpha_n$ denote the other roots. From Lemma 3.16 and for $i \leq k$, we see that $\max_{j>k} |\alpha_i - \alpha_j| \leq a^{\frac{n}{2(n-k)}}$ and thus $\prod_{j\neq i} |\alpha_i - \alpha_j| \leq \frac{2^{k-1}}{a^{n/2-n/k}}$. From (3.2), the latter inequality implies that the k first entries in the last row of $V_{g_{n,k,a}}^{-1}$ have magnitudes at least $\frac{a^{n/2-n/k}}{2^{k-1}}$. This gives us the claim.

Proposition 3.2 shows how to define polynomials for which the expansion factor is small and the inverse Vandermonde has very large entries. The following is an example. Note that there is some flexibility in the choice of a and k with respect to n to achieve the desired behavior. For example, one can also fix a and look for $k \leq C \log(n)$ for a constant C > 0.

Corollary 3.3. For k = 3 and $5 \le a \le \text{poly}(n)$, the polynomials $g_{n,3,a}$ satisfy

 $\mathsf{EF}(g_{n,3,a}) \le \mathsf{poly}(n)$ and $\|V_{q_{n,3,a}}^{-1}\|_{\mathsf{F}} \ge 2^{\Omega(n)}$.

We turn now to the second result of this section. We will make use of the observation that for any degree n polynomial f, we have that $\mathsf{EF}(f) \leq (2n-1) \cdot \max\{||x^k \mod f||_{\infty} : 0 \leq k \leq 2n-2\}.$

Lemma 3.17. Let $f(x) \in \mathbb{Z}[x]$ be a monic polynomial of degree n with distinct roots $\alpha_1, \ldots, \alpha_n$. Then $\mathsf{EF}(f) \leq (2n-1) \cdot ||V_f^{-1}||_{\mathsf{F}} \cdot ||V_f||_{\mathsf{F}}^2$.

Proof. Let C_f be the companion matrix of f, *i.e.*,

$$C_f = \begin{pmatrix} 0 & 0 & \dots & 0 & -f_0 \\ 1 & 0 & \dots & 0 & -f_1 \\ 0 & 1 & \dots & 0 & -f_2 \\ \vdots & & & & \\ 0 & 0 & \dots & 1 & -f_{n-1} \end{pmatrix} \in \mathbb{Z}^{n \times n}$$

Since the roots α_i of f are distinct, the matrix C_f is diagonalizable as follows

$$V_f \cdot C_f \cdot V_f^{-1} = \mathsf{diag}(\alpha_1, \dots, \alpha_n),$$

where $\operatorname{diag}(\alpha_1, \ldots, \alpha_n)$ is a diagonal $n \times n$ matrix having on the (i, i) position the root α_i of f. It means that for any k > 0, we have that $C_f^k = V_f^{-1} \cdot \operatorname{diag}(\alpha_1^k, \ldots, \alpha_n^k) \cdot V_f$. The last column of C_f corresponds to the coefficients of the polynomial $x^n \mod f$. By induction, it can be proved that actually for any $k \geq 1$, the last column of C_f^k corresponds to the coefficients of the polynomial $x^{n+k-1} \mod f$. As a consequence,

$$\begin{split} \mathsf{EF}(f) &\leq (2n-1) \cdot \max\{||C_f^k||_{\mathsf{F}} : 1 \leq k \leq n-1\} \\ &= (2n-1) \cdot \max\{||V_f^{-1} \cdot \operatorname{diag}(\alpha_1^k, \dots, \alpha_n^k) \cdot V_f||_{\mathsf{F}} : 1 \leq k \leq n-1\} \end{split}$$

The Frobenius norm is submultiplicative and for any $1 \le k \le n-1$, we have that

$$||\operatorname{diag}(\alpha_1^k,\ldots,\alpha_n^k)||_{\mathsf{F}} \leq ||V_f||_{\mathsf{F}},$$

which allow us to write $\mathsf{EF}(f) \leq (2n-1) \cdot ||V_f^{-1}||_{\mathsf{F}} \cdot ||V_f||_{\mathsf{F}}^2$ and to conclude.

CHAPTER 4

THE MIDDLE-PRODUCT LEARNING WITH ERRORS PROBLEM

In this chapter, we introduce the Middle-Product Learning With Errors problem (MP-LWE), a structured variant of LWE which makes use of the middle-product between two polynomials modulo an integer q. We show that MP-LWE is at least as hard as $PLWE^{(f)}$ for many polynomials f. The noise growth in the reduction is proportional to the expansion factor of f. We also show that MP-LWE remains hard even if the secrets are drawn from a distribution which produces small elements with high probability.

This chapter is mainly based on two articles: a joint work [RSSS17] with Amin Sakzad, Damien Stehlé and Ron Steinfeld accepted for publication at Crypto 2017 and another joint work [BDH+20] with Shi Bai, Dipayan Das, Ryo Hiromasa, Amin Sakzad, Damien Stehlé, Ron Steinfeld and Zhenfei Zhang which has been accepted at PKC 2020.

Contents

4.1	Introduction	64
4.2	Contributions	65
	4.2.1 Follow-up works	65
4.3	The middle-product of two polynomials	66
4.4	Middle-product learning with errors	67
4.5	Hardness of MP-LWE	68
4.6	A different way to see the PLWE to MP-LWE reduction $\ldots \ldots \ldots \ldots$	70
4.7	Hardness of MP-LWE with small secrets	72

4.1 Introduction

The PSIS and PLWE problems have been introduced as variants of SIS and LWE which led to more efficient cryptographic constructions. Still, while SIS and LWE enjoy reductions from some presumably hard worst-case lattice problems for general lattices, their corresponding algebraic variants PSIS and PLWE only enjoy reductions from the ApproxSVP problem restricted to the class of so-called *ideal lattices* ([LM06, PR06, SSTX09]). We refer to this variant of ApproxSVP as ApproxSVP^(f), since the ideal lattices depend on the polynomial f which parameterizes PSIS and PLWE. Concretely, the ideal lattices correspond to ideals of the ring $\mathbb{Z}[x]/f$.

The hardness of $\mathsf{ApproxSVP}^{(f)}$ has been investigated in a sequence of works. In the case of lattices corresponding to principal ideals in cyclotomic rings of prime-power conductor, Cramer et al. [CDPR16] gave a quantum polynomial time algorithm for solving the ApproxSVP^(f) problem with approximation factor $2^{O(\sqrt{n})}$ (where n is the degree of f) which works by first computing a generator of the ideal using [BS16] and then shortening it using the so-called *log-unit lattice*. Later on, Cramer et al. [CDW17] extended the algorithm from [CDPR16] to any ideal in a cyclotomic ring of prime-power conductor. Moreover, under some conditions on p, q, α and β , Holzer et al. [HWB17] extended [CDPR16] in the case of principal ideals in cyclotomic rings of conductor $p^{\alpha}q^{\beta}$. As a comparison, for such approximation factors and arbitrary lattices, the best known algorithms [Sch87] run in time $2^{O(\sqrt{n})}$. We mention that the development of faster algorithms for solving the ApproxSVP^(f) problem in cyclotomic number fields motivated the choice of non-cyclotomic polynomials in [BCLvV16]. Bauch et al. $[BBdV^+17]$ showed that in the case of multiquadratic fields, similar discrepencies between the hardness of ApproxSVP and Approx $SVP^{(f)}$ arise for principal ideals. Later on, building upon [CDW17], Pellet-Mary *et al.* [PHS19] improved the trade-offs for solving $\mathsf{ApproxSVP}^{(f)}$ in lattices corresponding to any ideal of the ring of integers of an arbitrary number field K, both classically and quantumly, at the cost of allowing the algorithm to perform some pre-computations. The pre-computations depend only on the ambient number field K and not on the lattice chosen and have run-time exponential in $\log(\Delta)$, where Δ is the discriminant of K. For instance, in the case of prime power cyclotomic fields, once the pre-processing is done, for a $2^{\widetilde{O}(n^{\alpha})}$ approximation factor, their algorithm takes $2^{\widetilde{O}(n^{1-2\alpha})} + \operatorname{poly}(n)$ time quantumly and $2^{\widetilde{O}(n^{1-2\alpha})} + 2^{\widetilde{O}(\sqrt{n})}$ classically.

One may wonder how hard $\mathsf{PSIS}^{(f)}$ or $\mathsf{PLWE}^{(f)}$ actually are, since a polynomial f for which the $\mathsf{ApproxSVP}^{(f)}$ problem is easy (or easier than $\mathsf{ApproxSVP}$) does not necessarily correspond to an easy $\mathsf{PSIS}^{(f)}$ or $\mathsf{PLWE}^{(f)}$ instance.

It could happen that $\mathsf{PSIS}^{(f)}$ or $\mathsf{PLWE}^{(f)}$ are easy to solve for some polynomials f, and hard for others. For instance, if f has a linear factor over the integers, then $\mathsf{PSIS}^{(f)}$ and $\mathsf{PLWE}^{(f)}$ are computationally easy (we note that the reductions from $\mathsf{ApproxSVP}^{(f)}$ require f to be irreducible). Apart from the very specific case of field extensions [GHPS12], hardness on K seems unrelated to hardness on another field K'. Finding weak f's for PLWE has been investigated in a series of articles [EHL14, ELOS15, CLS19, CLS16], but the respective attacks work only for error distributions with small width relative to the geometry of the corresponding ring [CIV16a, CIV16b, Pei16].

This lack of understanding of which f's correspond to hard $\mathsf{PLWE}^{(f)}$ problems motivates research into problems that are provably as hard as $\mathsf{PSIS}^{(f)}$ or $\mathsf{PLWE}^{(f)}$ for the hardest f in a large class of polynomials, while preserving the efficiency advantages of these problems. In [Lyu16], Lyubashevsky introduced a variant of $\mathsf{PSIS}^{(f)}$ which enjoys the above desirable property. This new problem, which we are going to call $\mathsf{PSIS}^{\emptyset}$, is not parametrized by a specific polynomial f, but only by the degree n. The main result in [Lyu16] is a reduction from $\mathsf{PSIS}^{(f)}$ to $\mathsf{PSIS}^{\emptyset}$ which works for all f's in a family of polynomials of size exponential in n. As a result, $\mathsf{PSIS}^{\emptyset}$ serves as an alternative cryptographic foundation that hedges against the risk that $\mathsf{PSIS}^{(f)}$ is easy to solve for some f as long as $\mathsf{PSIS}^{(f)}$ stays hard for some f in the family.

4.2 Contributions

Our first contribution is the introduction of an LWE counterpart to Lyubashevsky's $\mathsf{PSIS}^{\emptyset}$ problem: the middle-product learning with errors problem (MP-LWE). This problem is defined in Section 4.4 using the so-called middle-product of two polynomials. Let $n, q \geq 2$. We let $\mathbb{Z}_q^{\leq n}[x]$ denote the set of polynomials with coefficients in \mathbb{Z}_q and degree < n. For $a \in \mathbb{Z}_q^{\leq n}[x]$ and $s \in \mathbb{Z}_q^{\leq 2n-1}[x]$, we let $a \odot_n s = \lfloor (a \cdot s \mod x^{2n-1})/x^{n-1} \rfloor \in \mathbb{Z}_q^{\leq n}[x]$ denote the polynomial obtained by multiplying a and s and keeping only the middle n coefficients. The MP-LWE problem with parameters $n, q \geq 2$ and $\alpha \in (0, 1)$, consists in distinguishing arbitrarily many samples (a_i, b_i) uniform in $\mathbb{Z}_q^{\leq n}[x] \times (\mathbb{R}/q\mathbb{Z})^{\leq n}[x]$, from the same number of samples (a_i, b_i) with a_i uniform in $\mathbb{Z}_q^{\leq n}[x]$ and $b_i = a_i \odot_n s + e_i$, where each coefficient of e_i is sampled from the Gaussian distribution of standard deviation $\alpha \cdot q$, and s is uniformly chosen in $\mathbb{Z}_q^{\leq 2n-1}[x]$. The reversed coefficient vector of the middle-product of two polynomials is in fact equal to the product of the Toeplitz matrix associated to one polynomial by the reversed coefficient vector of the second polynomial.

In Section 4.5 we give a reduction from (decision) $\mathsf{PLWE}^{(f)}$ to (decision) $\mathsf{MP-LWE}$ of parameter n, for every monic f of degree n whose constant coefficient is coprime with q. We prove this result in two steps. First, we map the PLWE samples to a variant of $\mathsf{MP-LWE}$ whose error distribution depends on the matrix \mathbf{M}_f parameterized by the polynomial f and whose corresponding secret is non-uniform. In the second step, we re-randomize the secret and remove the dependency on f of the error by adding a compensating Gaussian distribution. In the end, the noise parameter of the $\mathsf{MP-LWE}$ samples amplifies linearly with the expansion factor of f and can for example be set to handle all monic polynomials $f = x^n + g$ with constant coefficient coprime with q, deg $g \leq n/2$ and $||g|| \leq n^c$ for an arbitrary c > 0. For any c, this set of f's has exponential size in n. We note that similar restrictions involving the expansion factor appeared before in [LM06, SSTX09].

If the free coefficient of the polynomial f(x) is invertible in \mathbb{Z}_q , f(x) is invertible in the ring $\mathbb{Z}_q[[x]]$ and the expansion of $\frac{1}{f(x)}$ as a formal series is closely related to the matrix \mathbf{M}_f . This observation allows us to rewrite the reduction from PLWE to MP-LWE in a more algebraic way in Section 4.6.

Finally, in Section 4.7, we show that the reduction from $\mathsf{PLWE}^{(f)}$ to $\mathsf{MP-LWE}$ still holds if the errors are drawn from a discrete distribution and the secrets are sampled from a distribution which produces small elements with high probability. Discretizing the noise distribution is more convenient in real applications and can be achieved via routine techniques. Oppositely, having the secret distribution take small values compared to q is not straightforward. In contrast with the result from Section 4.5, this reduction requires a condition on the noise parameter α which arrises when we approximate the distribution of the sum of two random discrete variables by a new discrete distribution as in Lemma 2.7. The condition on α is implied by a lower bound on the smallest singular value of the matix \mathbf{M}_f and we manage to bound from below the smallest singular value of the matrix \mathbf{M}_f for an exponentially large family of polynomials f.

4.2.1 Follow-up works

In this section, we briefly describe the connection of some works built upon MP-LWE with the results presented in this chapter.

4.2.1.1 Middle-Product Learning With Rounding

At Asiacrypt 2019, Bai *et al.* [BBD⁺19] introduced the Middle-Product Computational Learning With Rounding (MP-CLWR) problem, as a natural adaption to the middle-product context of the computational Learning With Rounding (LWR) problem over rings [CZZ18]. The motivations of this new problem were twofold. Firstly, on the security front, MP-CLWR is at least as hard as MP-LWE, and it thus relies on the hardness of solving certain lattice problems on a large class of lattices. Secondly, by using rounding, the MP-CLWR problem avoids the Gaussian error sampling which could be costly and easily exploitable by an attacker. To illustrate the cryptographic relevance of MP-CLWR, the authors of [BBD⁺19] build a public key encryption scheme IND-CPA secure in the random oracle model under the MP-CLWR hardness assumption, with the same asymptotic efficiency as the one based on MP-LWE that we present in Chapter 5.

4.2.1.2 A general framework for all LWE variants

Peikert and Pepin ([PP19]) proposed at TCC 2019 a general framework to analyse all the (algebraic) variants of LWE in the literature.

In particular, they define a new problem \mathcal{L} -LWE parameterized by a lattice \mathcal{L} in a number field Kand show that PLWE, primal/dual-RLWE, Order-LWE and Module-LWE can be obtained as particular cases of this problem instantiated with suitable lattices \mathcal{L} . Then they show that for any two lattices $\mathcal{L}' \subseteq \mathcal{L}$ of a number field K, under some conditions on their so-called *coefficient rings* and $|\mathcal{L} : \mathcal{L}'|$, there is an error preserving reduction from \mathcal{L} -LWE to \mathcal{L}^{\vee} -LWE. Using the above result, they obtain the hardness of a variant of PLWE with secret belonging to $\mathbb{Z}[\alpha]^{\vee}$ instead of $\mathbb{Z}[\alpha]$, which they call dual-PLWE, by directly reducing dual-RLWE to this new problem. Moreover, they rewrite the definition of the MP-LWE problem in terms of biliniar maps and use it to give a reduction from dual-PLWE to MP-LWE. As a consequence, they reobtain the hardness of MP-LWE from Section 4.4 based on the hardness of solving certain lattice problems in a wide class of lattices.

4.3 The middle-product of two polynomials

In this section, we recall the definition of the middle-product of two polynomials and we exhibit its relationship with Toeplitz matrices. Let R be a ring. Assume we have two polynomials a and b of degrees $\langle d_a \rangle$ and $\langle d_b \rangle$, respectively, and $d_a + d_b - 1 = d + 2k$ for some integers d and k. Then the middle-product of size d of a and b is obtained by multiplying a and b, deleting the left coefficients of $1, x, \ldots, x^{k-1}$, deleting the right coefficients of $x^{k+d}, x^{k+d+1}, \ldots, x^{d+2k-1}$ and dividing what remains in the middle by x^k .

Definition 4.1. Let d_a, d_b, d, k be integers such that $d_a + d_b - 1 = d + 2k$. The middle-product $\odot_d : R^{\leq d_a}[x] \times R^{\leq d_b}[x] \to R^{\leq d}[x]$ is the map:

$$(a,b) \mapsto a \odot_d b = c_k + c_{k+1} \cdot x + \dots + c_{k+d-1} \cdot x^{d-1},$$

where $a(x) \cdot b(x) = \sum_{i=0}^{d+2k-1} c_i \cdot x^i$. We use the same notation \odot_d for every d_a, d_b such that $d_a + d_b - 1 - d$ is non-negative and even.

The middle-product of polynomials is used in computer algebra to accelerate computations in polynomial rings (see, e.g., [Sho99, HQZ04]). As it is part of the output of polynomial multiplication, it can be computed with a number of ring additions and multiplications that is quasi-linear number in $d_a + d_b$, but faster algorithms also exist [HQZ04].

Lemma 4.1. Let d, k > 0. Let $r \in R^{<k+1}[x]$ and $a \in R^{<k+d}[x]$ and $b = r \odot_d a$. Then $\overline{\mathbf{b}} = \operatorname{Toep}^{d,k+1}(r) \cdot \overline{\mathbf{a}}$. In other words, we have $\mathbf{b} = \operatorname{Toep}^{d,k+1}(r) \cdot \overline{\mathbf{a}}$.

Proof. We first note that $\mathsf{Toep}^{d,2k+d}(r \cdot a) = \mathsf{Toep}^{d,k+1}(r) \cdot \mathsf{Toep}^{k+d,k+d}(a)$. Thus, by definition of the middle-product, we have that the coefficients of b appear in the first row of $\mathsf{Toep}(r \cdot a)$, namely $b_i = \mathsf{Toep}^{d,2k+d}(r \cdot a)_{1,k+i+1}$ for i < d. But since $\mathsf{Toep}(r \cdot a)$ is constant along its diagonals, we

also have that b appear (in reversed order) in the (k + d)-th column of $\mathsf{Toep}^{d,2k+d}(r \cdot a)$, namely $b_i = \mathsf{Toep}^{d,2k+d}(r \cdot a)_{d-i,k+d}$ for i < d. Therefore, vector $\overline{\mathbf{b}}$ is the (k+d)-th column of $\mathsf{Toep}^{d,2k+d}(r \cdot a)$, which is equal to $\mathsf{Toep}^{d,k+1}(r) \cdot \mathbf{a}'$, where \mathbf{a}' is the (k+d)-th column of $\mathsf{Toep}^{k+d,k+d}(a)$. Since $\mathsf{Toep}^{k+d,k+d}(a)$ is constant along its diagonals, its first row is equal to its reversed (k+d)-th column, so $\mathbf{a}' = \overline{\mathbf{a}}$, as required.

The above lemma and the example following Lemma 2.2 indicate that the multiplication modulo $f = x^n + 1$ can be easily converted into a middle-product. We will further exploit this connection in Section 4.5 using Lemma 2.1 without restricting the choice of the polynomial f.

Observation 4.3.1. Notice that using the notations from the previous lemma, we also have the following equality:

											$\left(a_{k+d-1}\right)$
	10			0					١		a_{k+d-2}
		•••	•••	0	r_0	r_1	•••	r_{k-1}	r_k		a_{k+d-3}
	0	•••	•••	r_0	r_1	r_2	•••	r_k	0		:
$\mathbf{b} =$	0		· · ·	· · ·	· • *	· · ·	· · ·	0	0	•	:
	0	r_0	[.]	· · ·	·	r_k	0		0		:
	r_0	r_1			r_k	0			0 /		:
											$\langle a_0 \rangle$

The middle-product is an additive homomorphism when either of its inputs is fixed. As a consequence of the associativity of matrix multiplication and Lemma 4.1, the middle-product satisfies the following "associativity" property.

Lemma 4.2. Let d, k, n > 0. For all $r \in R^{<k+1}[x]$, $a \in R^{<n}[x]$, $s \in R^{<n+d+k-1}[x]$, we have that $r \odot_d (a \odot_{d+k} s) = (r \cdot a) \odot_d s$.

Proof. Note first that the degree bounds match. Now, by Lemma 4.1, the vector associated to the reverse of $r \odot_d (a \odot_{d+k} s)$ is $\mathsf{Toep}^{d,k+1}(r) \cdot (\mathsf{Toep}^{d+k,n}(a) \cdot \overline{s})$. Similarly, the vector associated to the reverse of $(r \cdot a) \odot_d s$ is $\mathsf{Toep}^{d,k+n}(r \cdot a) \cdot \overline{s}$. The result follows from observing that $\mathsf{Toep}^{d,k+1}(r) \cdot \mathsf{Toep}^{d+k,n}(a) = \mathsf{Toep}^{d,k+n}(r \cdot a)$.

4.4 Middle-product learning with errors

In this section, we use the middle-product introduced in the previous section to define a new computational problem: Middle-Product Learning With Errors (MP-LWE). Before stating the MP-LWE problem, we first introduce the distribution its definition relies on.

Definition 4.2 (MP-LWE distribution). Let $n, d > 0, q \ge 2$, and χ a distribution over $\mathbb{R}_q^{< d}[x]$. For $s \in \mathbb{Z}_q^{< n+d-1}[x]$, we define the distribution $\mathsf{MP}_{q,n,d,\chi}(s)$ over $\mathbb{Z}_q^{< n}[x] \times \mathbb{R}_q^{< d}[x]$ as the one obtained by: sampling $a \leftrightarrow U(\mathbb{Z}_q^{< n}[x]), e \leftrightarrow \chi$ and returning $(a, b = a \odot_d s + e)$.

Definition 4.3 (MP-LWE). Let n, d > 0. Let χ_1 and χ_2 be distributions over $\mathbb{R}_q^{<d}[x]$ and $\mathbb{Z}_q^{<n+d-1}[x]$, respectively. The decision MP-LWE_{q,n,d,χ_1,χ_2} problem consists in distinguishing between arbitrarily many samples from MP_{q,n,d,χ_1}(s) and the same number of uniform samples in $\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x]$, with non-negligible probability over the choice of $s \leftrightarrow \chi_2$.

We can also define a search variant of $\mathsf{MP}\text{-}\mathsf{LWE}_{q,n,d,\chi_1,\chi_2}$, which consists in computing $s \in \mathbb{Z}_q^{< n+d-1}[x]$ from arbitrarily many samples from the distribution $\mathsf{MP}_{q,n,d,\chi_1}(s)$, where s has been sampled from the χ_2 distribution. When χ_2 is the uniform distribution on $\mathbb{Z}_q^{< n+d-1}[x]$, we simply write $\mathsf{MP}\text{-}\mathsf{LWE}_{q,n,d,\chi_1}$ instead of $\mathsf{MP}\text{-}\mathsf{LWE}_{q,n,d,\chi_1,\chi_2}$.

The decision/search MP-LWE_{q,n,d, χ_1,χ_2} problems can be viewed as variants of the corresponding LWE problem in which the samples are correlated. For instance, thanks to Lemma 4.1, the decision variant of MP-LWE_{q,n,d, χ_1,χ_2} can be restated as follows: given many samples $(\mathsf{Toep}^{d,n}(a_i), \overline{\mathbf{b}}_i) \in \mathbb{Z}_q^{d \times (n+d-1)} \times \mathbb{R}_q^d$ for uniformly chosen $a_i \in \mathbb{Z}_q^{< n}[x]$, decide if the vectors $\overline{\mathbf{b}}_i$ are uniformly sampled in \mathbb{R}_q^d or are of the form $\overline{\mathbf{b}}_i = \mathsf{Toep}^{d,n}(a_i) \cdot \overline{\mathbf{s}} + \overline{\mathbf{e}}_i$ for some $e_i \leftrightarrow \chi_1$ and some common secret $s \leftrightarrow \chi_2$.

Interestingly, Toeplitz matrices have already been used in cryptography in the context of symmetric key authentication protocols. In [GRS08], Gilbert *et al.* propose the Random-HB[#] and HB[#] protocols which improve the HB⁺ protocol [JW05] of Juels and Weis in terms of security and practicality. The security of HB[#] relies on the conjectured hardness of solving the so-called "Toeplitz-MHB" problem. The Toeplitz-MHB problem makes use of a binary Toeplitz matrix and binary vectors. We omit here its formal definition which can be found in [GRS08]. One could think to extend this definition by removing the binary condition put on the objects used in the following way.

Definition 4.4 (LWE-Toeplitz-MHB_{q,n,d,\chi}). Let $n, d > 0, q \ge 2$ and χ a distribution over \mathbb{R}^d . Let S be a random secret Toeplitz matrix in $\mathbb{Z}_q^{n \times d}$. Given many samples $(a_i, a_i S + e_i)$, where $a_i \leftrightarrow U(\mathbb{Z}_q^n)$ and $e_i \leftrightarrow \chi$, and a vector $a \in \mathbb{Z}_q^n$ uniformly chosen, find aS.

Since the matrix S is Toeplitz, it is uniquely defined by its first column and its first row. Suppose its first column is $(s_{n-1}, \ldots, s_0)^t$ and its first row is $(s_{n-1}, \ldots, s_{n+d-2})$. Notice that for any vector $a_i \in \mathbb{Z}_q^n$, a_i S is the vector of coefficients of the polynomial $a_i \odot_d s$, where $s := \sum_{i=0}^{n+d-2} s_i x^i$ and a_i is naturally identified with a polynomial in $\mathbb{Z}_q^{< n}[x]$. As a consequence, the MP-LWE_{q,n,d, χ} problem and the LWE-Toeplitz-MHB_{q,n,d, χ} problem resemble in terms of their inputs.

The LWE-Toeplitz-MHB_{q,n,d, χ} problem trivially reduces to MP-LWE_{q,n,d, χ}. While Toeplitz-MHB is only conjectured to be hard in [GRS08], in the next section we give concrete evidence of the hardness of MP-LWE based on the hardness of solving the PLWE^f problem for many polynomials f.

4.5 Hardness of MP-LWE

In this section, we give a reduction from $\mathsf{PLWE}^{(f)}$ to $\mathsf{MP-LWE}$ which works for many polynomials f. We manage to get better parameters compared to those in Theorem 3.6 from [RSSS17] by making use of Lemma 2.4, which improves on Lemma 2.8 from [RSSS17].

Theorem 4.1. Let n, d > 0, $q \ge 2$, and $\alpha \in (0, 1)$. For S > 0, we let $\mathcal{F}(S, d, n)$ denote the set of polynomials $f \in \mathbb{Z}[x]$ that are monic, have constant coefficient coprime with q, have degree m in [d, n] and that satisfy $\mathsf{EF}(f) < S$. Then there exists a ppt reduction from $\mathsf{PLWE}_{q,D_{\alpha \cdot q}}^{(f)}$ for any $f \in \mathcal{F}(S, d, n)$ to $\mathsf{MP-LWE}_{q,n,d,D_{\alpha',q}}$ with $\alpha' = \alpha \sqrt{dS}$.

Proof. We first reduce $\mathsf{PLWE}^{(f)}$ to a variant of $\mathsf{MP-LWE}$ whose only dependency on f lies in the noise distribution (see Lemma 4.3 below). Then we remove the latter dependency, by adding a compensating Gaussian distribution (see Lemma 4.4 below). The bound on the magnitude of matrix \mathbf{M}_f from Lemma 2.4 for $\chi = D_{\alpha \cdot q}$ implies that

$$\|\boldsymbol{\Sigma}_0\| = \alpha^2 q^2 \|\mathbf{J} \cdot \mathbf{M}_f^d\|^2 = \alpha^2 q^2 \|\mathbf{M}_f^d\|^2 \le d(\alpha q \mathsf{EF}(f))^2 < d(\alpha q S)^2.$$

Hence, taking $\alpha' q = \alpha q \sqrt{dS}$ completes the proof.

Lemma 4.3. Let $n, d > 0, q \ge 2$, and χ a distribution over $\mathbb{R}^{\leq d}[x]$. Then there exists a ppt reduction from $\mathsf{PLWE}_{q,\chi}^{(f)}$ for any monic $f \in \mathbb{Z}[x]$ with constant coefficient coprime with q and degree $m \in [d, n]$, to $\mathsf{MP-LWE}_{q,n,d,\mathbf{J}\cdot\mathbf{M}_{f}^{d}\cdot\chi}$. Here, matrix \mathbf{M}_{f}^{d} is the one obtained by keeping only the first d rows of \mathbf{M}_{f} , and $\mathbf{J} \in \mathbb{Z}^{d \times d}$ is the one with 1's on the anti-diagonal and 0's everywhere else.

Proof. We describe below an efficient randomized mapping ϕ that takes as input a pair $(a_i, b_i) \in \mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$ and maps it to a pair $(a'_i, b'_i) \in \mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x]$, such that ϕ maps $U(\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f)$ to $U(\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x])$ and $\mathsf{P}_{q,\chi}^{(f)}(s)$ to $\mathsf{MP}_{q,n,d,\chi'}(s')$, for some s' that depends on s and some χ' that depends on χ and f.

The reduction is then as follows:

- Sample $t \leftarrow U(\mathbb{Z}_a^{< n+d-1}[x])$.
- Each time the MP-LWE oracle requests a new sample, ask for a fresh PLWE sample (a_i, b_i) , compute $(a'_i, b'_i) = \phi(a_i, b_i)$ and give $(a'_i, b'_i) + (0, a'_i \odot_d t)$ to the MP-LWE oracle.
- When MP-LWE terminates, return its output.

Assuming ϕ satisfies the specifications above, the reduction maps uniform samples to uniform samples, and $\mathsf{P}_{q,\chi}^{(f)}(s)$ samples for a uniform s that is common to all samples to $\mathsf{MP}_{q,n,d,\mathbf{J}\cdot\mathbf{M}_{f}^{d}\cdot\chi}(s'+t)$ samples for a uniform s' + t that is common to all samples.

We now describe ϕ . Let $(a_i, b_i) \in \mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$ be an input pair. Let m denote the degree of f. We sample $r_i \leftarrow U(\mathbb{Z}_q^{\leq n-m}[x])$ and set $\phi(a_i, b_i) = (a'_i, b'_i)$ with:

$$a'_i = a_i + f \cdot r_i \in \mathbb{Z}_q^{< n}[x], \quad \overline{\mathbf{b}'_i} = \mathbf{M}_f^d \cdot \mathbf{b}_i \in \mathbb{R}_q^{< d}[x].$$

As a_i and r_i are uniformly distributed in $\mathbb{Z}_q^{\leq m}[x]$ and $\mathbb{Z}_q^{\leq n-m}[x]$ respectively, the polynomial a'_i is uniformly distributed in $\mathbb{Z}_q^{\leq n}[x]$ (we refer to [Lyu16, Lemma 2.10] for a fully detailed proof).

Further, if b_i is uniformly distributed, then so is its coefficient vector \mathbf{b}_i , and so is $\mathbf{M}_f^d \cdot \mathbf{b}_i$. Indeed, as the constant coefficient is coprime with q, the matrix \mathbf{M}_f is invertible modulo q (reordering its columns makes it triangular, with diagonal coefficients invertible modulo q).

Now, assume that $b_i = a_i \cdot s + e_i$, for some $s \in \mathbb{Z}_q[x]/f$ and $e_i \leftrightarrow \chi$. Thanks to Subsection 2.2, we know that $\operatorname{Rot}_f(b_i) = \operatorname{Rot}_f(a_i) \cdot \operatorname{Rot}_f(s) + \operatorname{Rot}_f(e_i)$, and, by taking the first columns and d first rows, we have

$$\begin{split} \mathbf{M}_{f}^{d} \cdot \mathbf{b}_{i} &= \operatorname{Rot}_{f}^{d}(a_{i}) \cdot \mathbf{M}_{f} \cdot \mathbf{s} + \mathbf{M}_{f}^{d} \cdot \mathbf{e}_{i} \\ &= \operatorname{Rot}_{f}^{d}(a_{i}') \cdot \mathbf{M}_{f} \cdot \mathbf{s} + \mathbf{M}_{f}^{d} \cdot \mathbf{e}_{i} \\ &= \operatorname{Toep}^{d,n}(a_{i}') \cdot \operatorname{Rot}_{f}^{d+n-1}(1) \cdot \mathbf{M}_{f} \cdot \mathbf{s} + \mathbf{M}_{f}^{d} \cdot \mathbf{e}_{i} \\ &= \operatorname{Toep}^{d,n}(a_{i}') \cdot \overline{\mathbf{s}'} + \mathbf{M}_{f}^{d} \cdot \mathbf{e}_{i}, \end{split}$$

where $\mathbf{s}' = \overline{\mathsf{Rot}_f^{d+n-1}(1) \cdot \mathbf{M}_f \cdot \mathbf{s}}$. Since $\mathbf{b}'_i = \overline{\mathbf{M}_f^d \cdot \mathbf{b}_i} = \overline{\mathsf{Toep}(a'_i) \cdot \mathbf{s}'} + \overline{\mathbf{M}_f^d \cdot \mathbf{e}_i}$, we get that $\mathbf{e}'_i = \overline{\mathbf{M}_f^d \cdot \mathbf{e}_i}$, which makes the distribution in MP-LWE equal to the claimed $\mathbf{J} \cdot \mathbf{M}_f^d \cdot \chi$. This completes the proof. \Box

We now remove the dependency on f of the noise distribution.

Lemma 4.4. Let $n, d > 0, q \ge 2$. Let $\sigma' > 0$. Let $\Sigma_0 \in \mathbb{R}^{d \times d}$ be symmetric definite positive matrix with $\|\Sigma_0\| < (\sigma')^2$. Then there exists a ppt reduction from MP-LWE_{q,n,d,D_{$\Sigma_0}} to MP-LWE_{q,n,d,D_{<math>\sigma'}}.</sub></sub>$ </sub></sub>

Proof. The reduction is as follows. We first note that, there exists a positive definite matrix Σ' , such that $\Sigma_0 + \Sigma' = (\sigma')^2 \cdot \mathsf{Id}_d$. The positive definiteness is guaranteed by fact that $\|\Sigma_0\| < (\sigma')^2$. Then, for any MP-LWE_{q,n,d,D_{$\Sigma_0}} input sample <math>(a_i, b_i)$, we sample $e'_i \leftrightarrow D_{\Sigma'}$ and compute $(a'_i, b'_i) = (a_i, b_i + e'_i)$.</sub></sub>

Observe that the reduction maps uniform samples to uniform samples, and $\mathsf{MP}_{q,n,d,D_{\Sigma_0}}(s)$ samples to $\mathsf{MP}_{q,n,d,D_{\sigma'}}(s)$ samples. This completes the proof.

Notice that we do not use the monocity of the polynomial f in the reduction and the only condition on the coefficients of f (*i.e.* $(f_0, q) = 1$) is necessary to preserve the uniformity by multiplication with the matrix \mathbf{M}_f . Still, we keep this monocity condition in the statement of the theorem in order to be consistent with the definition of the PLWE problem.

4.6 A different way to see the PLWE to MP-LWE reduction

In this section, we first embed the ring $\mathbb{Z}_q^{< n}[x]$ into $\mathbb{Z}_q[[x]]$, naturally identifying a polynomial $a \in \mathbb{Z}_q^{< n}[x]$ with the formal series whose first n coefficients are equal to the coefficients of a and whose other coefficients are set to 0, and then rewrite the definition of the middle-product of a and b by allowing the element b to be a formal series instead of a polynomial. This allows us to give an algebraic interpretation of the reduction from Theorem 4.1. We use the notations introduced in Section 2.2.

Definition 4.5. The middle-product $\odot_d : \mathbb{Z}_q^{< n}[x] \times \mathbb{Z}_q[[x]] \to \mathbb{Z}_q^{< d}[x]$ is the map:

$$(a,b) \mapsto a \odot_{d} b = [ab]_{n-1}^{n+d-1}$$

The multiplication $a \cdot b$ is done in the ring $\mathbb{Z}_q[[x]]$ by identifying the polynomial $a \in \mathbb{Z}_q^{<n}[x]$ with the formal series whose first n coefficients are equal to the coefficients of a and whose other coefficients are all set to 0. Notice that for a fixed $a \in \mathbb{Z}_q^{<n}[x]$ and any two formal series b and b' which coincide on their first n + d - 1 coefficients, we have that $a \odot_d b = a \odot_d b'$. This means that the following two definitions are just a restatement of the middle-product distribution and problem.

Definition 4.6 (MP-LWE distribution). Let $n, d > 0, q \ge 2$ and a distribution χ over $\mathbb{R}_q^{\leq d}[x]$. For $s \in \mathbb{Z}_q[[x]]$, we define the distribution $\mathsf{MP}_{q,n,d,\chi}(s)$ over $\mathbb{Z}_q^{\leq n}[x] \times \mathbb{R}_q^{\leq d}[x]$ as the one obtained by: sampling $a \leftrightarrow U(\mathbb{Z}_q^{\leq n}[x]), e \leftrightarrow \chi$ and returning $(a, b = a \odot_d s + e)$.

Definition 4.7 (MP-LWE). Let $n, d > 0, q \ge 2$. Let χ_1 and χ_2 distributions over $\mathbb{R}_q^{<d}[x]$ and $\mathbb{Z}_q[[x]]$, respectively. The decision MP-LWE_{n,d,q,χ_1} consists in distinguishing between arbitrarily many samples from $\mathsf{MP}_{q,n,d,\chi_1}(s)$ and the same number of samples from $U(\mathbb{Z}_q^{< n}[x] \times \mathbb{R}_q^{<d}[x])$, with non-negligible probability over the choices of $s \leftrightarrow \chi_2$.

Theorem 4.2. Let n, d, S > 0, $q \ge 2$ and $\alpha \in (0, 1)$. For any monic polynomial $f \in \mathbb{Z}[x]$ of degree $m \in [d, n]$ such that $(f_0, q) = 1$ and $\mathsf{EF}(f) < S$, there is a polynomial time reduction from $\mathsf{PLWE}_{q, D_{\alpha q}}^{(f)}$ to $\mathsf{MP-LWE}_{q, n, d, D_{\alpha' q}}$, where $\alpha' = \alpha \sqrt{dS}$.

Proof. It is enough to describe an efficient randomized mapping ϕ which takes as input a pair $(a, b) \in \mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$ and maps it to a pair $(a', b') \in \mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x]$ such that ϕ maps $U(\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f)$ to $U(\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x])$ and $\mathsf{P}_{q,\mathcal{D}_{\alpha q}}^{(f)}(s)$ to $\mathsf{MP}_{q,n,d,\mathcal{D}_{\alpha' q}}(s')$. First, we choose $s^* \leftrightarrow U(\mathbb{Z}_q[[x]])$. For each sample asked by the MP-LWE oracle, we ask the PLWE oracle for a sample (a_i, b_i) . Now we

For each sample asked by the MP-LWE oracle, we ask the PLWE oracle for a sample (a_i, b_i) . Now we choose $r_i \leftrightarrow U(\mathbb{Z}_q^{< n-m}[x])$ and set $a'_i(x) = \overline{a_i} + fr_i(x) = \overline{a_i}(x) \cdot x^{n-m} + \overline{f}(x) \cdot \overline{r_i}(x)$, which is uniform in $\mathbb{Z}_q^{< n}[x]$, using the same argument as in Lemma 4.3. Since the leading coefficient of f is invertible in \mathbb{Z}_q , \overline{f} is invertible in the ring of formal series and we can set

$$b'_i(x) = \left[\frac{\overline{b_i}(x)}{\overline{f}(x)}\right]_{m-1}^{m-1+d-1} + [a'_i(x)s^*(x)]_{n-1}^{n-1+d-1} + e'_i(x)$$

for some e'_i chosen from an error compensating Gaussian distribution.
We have two cases, depending on the distribution from which the sample (a, b) has been chosen.

If for any *i*, we have $b_i = a_i s + e_i \mod f$ for some common secret *s* and $e_i \leftrightarrow D_{\alpha q}$, then there exists a polynomial $k_i \in \mathbb{Z}_q[x]$ of degree $\leq m-2$ such that $b_i = a_i s + e_i - k_i f$ and $\overline{b_i}(x) = \overline{a_i}(x)\overline{s}(x) + \overline{e_i}(x) \cdot x^{m-1} - \overline{k_i}(x) \cdot \overline{f}(x)$.

Notice that

$$\frac{\overline{b_i}(x)}{\overline{f}(x)} = \frac{a_i'(x) - \overline{f}(x)\overline{r_i}(x)}{x^{n-m}} \cdot \frac{\overline{s}(x)}{\overline{f}(x)} + \frac{\overline{e_i}(x)}{\overline{f}(x)} \cdot x^{m-1} - \overline{k_i}(x)$$

$$= a_i'(x) \cdot \frac{\overline{s}(x)}{\overline{f}(x)} \cdot \frac{1}{x^{n-m}} + \frac{\overline{e_i}(x) \cdot x^{m-1}}{\overline{f}(x)} - \overline{k_i}(x) - \frac{\overline{r_i}(x)\overline{s}(x)}{x^{n-m}}$$

and

$$\left[\frac{\overline{b_i}(x)}{\overline{f}(x)}\right]_{m-1}^{m-1+d-1} = \left[a_i'(x) \cdot \frac{\overline{s}(x)}{\overline{f}(x)} + \frac{\overline{e_i}(x) \cdot x^{n-1}}{\overline{f}(x)} - \overline{k_i}(x) \cdot x^{n-m} - \overline{r_i}(x)\overline{s}(x)\right]_{n-1}^{n-1+d-1}$$

The degree of the polynomials $\overline{k_i}(x) \cdot x^{n-m}$ and $\overline{r_i}(x)\overline{s}(x)$ is less than n-1, which implies that

$$\left[\frac{\overline{b_i}(x)}{\overline{f}(x)}\right]_{m-1}^{m-1+d-1} = \left[a_i'(x) \cdot \frac{\overline{s}(x)}{\overline{f}(x)} + \frac{\overline{e_i}(x) \cdot x^{n-1}}{\overline{f}(x)}\right]_{n-1}^{n-1+d-1}$$

It follows that

$$b'_{i}(x) = \left[a'_{i}(x) \cdot \left(\frac{\overline{s}(x)}{\overline{f}(x)} + s^{*}(x)\right)\right]_{n-1}^{n-1+d-1} + \left[\frac{\overline{e_{i}}(x)}{\overline{f}(x)}\right]_{0}^{d-1} + e'_{i}(x).$$

In the following, we analyse the distribution of $\left[\frac{\overline{e_i}(x)}{\overline{f}(x)}\right]_0^{m-1}$. By writing $\frac{1}{\overline{f}(x)} = \sum_{j=0}^{\infty} z_j \cdot x^j$, we notice that the vector of the first *m* coefficients of $\frac{\overline{e_i}(x)}{\overline{f}(x)}$ is $\mathbb{Z}e_i$, where

By identifying the coefficients of $\frac{1}{\overline{f}(x)}$, we obtain that $z_0 = 1$ and

$$z_{j} = \begin{vmatrix} -f_{m-1} & -f_{m-2} & \cdots & -f_{m-j+1} & -f_{m-j} \\ -1 & -f_{m-1} & \cdots & -f_{m-j+2} & -f_{m-j+1} \\ 0 & -1 & \cdots & \cdots & \cdots \\ 0 & \cdots & \cdots & \cdots & \cdots \\ 0 & \cdots & \cdots & -1 & -f_{m-1} & -f_{m-2} \\ 0 & \cdots & \cdots & 0 & -1 & -f_{m-1} \end{vmatrix}$$

for any $j \in \{1, m-1\}$. One can now notice that $z_j = -[x^{m+j} \mod f]_0$ for any $j \in \{0, \ldots, m-1\}$, which means that the matrix Z is closely related to the matrix \mathbf{M}_f that we used in the proof of Theorem 4.1. Moreover, using the definition of the expansion factor and the fact that $z_{m-1} = 0$, we

get that $|z_j| < \mathsf{EF}(f)$ for any $j \le m - 1$.

We can see now that the vector of coefficients of the polynomial $\left[\frac{\overline{e_i}(x)}{\overline{f}(x)}\right]_0^{d-1}$ can be written as $\mathsf{Z}^{(d)} \cdot e_i$, where $\mathsf{Z}^{(d)} \in \mathbb{Z}_q^{d \times m}$ is the submatrix of Z containing only the first d rows. Using the same argument as in the proof of Lemma 2.4, we have that $||\mathsf{Z}^{(d)}||^2 \leq (\sqrt{d} \cdot \mathsf{EF}(f))^2 < (\sqrt{d} \cdot S)^2$. This implies that by adding a compensating Gaussian $e'_i(x)$ as in Theorem 4.1, the distribution of (a'_i, b'_i) will be exactly $\mathsf{MP}_{q,n,d,\mathcal{D}_{\alpha' q}}$.

If for any *i*, the element b_i was uniform in $\frac{\mathbb{R}_q[x]}{[f]}$, then $\left[\frac{\overline{b_i}(x)}{\overline{f}(x)}\right]_{m-1}^{m-1+d-1}$ will also be uniform in $\mathbb{R}_q^{< d}[x]$. Indeed, notice that $\left[\frac{\overline{b_i}}{\overline{f}}\right]_k = \frac{1}{[f]_m} \left([\overline{b_i}]_k - \sum_{j=1}^k [\overline{f}]_j \cdot \left[\frac{\overline{b_i}}{\overline{f}}\right]_{k-j}\right)$ for any $k \ge 0$. We set by definition $\left[\frac{\overline{b_i}}{\overline{f}}\right]_{-1} := -\frac{[b_i]_0}{[f]_0}$. Now we can see that

$$\left(\left[\frac{\overline{b_i}}{\overline{f}}\right]_{m-1},\ldots,\left[\frac{\overline{b_i}}{\overline{f}}\right]_{m-1+m-1}\right) = -\frac{1}{[f]_m} \left(\left[\frac{\overline{b_i}}{\overline{f}}\right]_{-1},\ldots,\left[\frac{\overline{b_i}}{\overline{f}}\right]_{m-2}\right) \cdot \mathsf{F}^m,$$

where F is the matrix

$$\mathsf{F} = \begin{pmatrix} 0 & 0 & \dots & 0 & [f]_0 \\ 1 & 0 & \dots & 0 & [f]_1 \\ 0 & 1 & \dots & 0 & [f]_2 \\ \vdots & & & & \\ 0 & 0 & \dots & 1 & [f]_{m-1} \end{pmatrix} \in \mathbb{Z}^{m \times m}.$$

Since b is uniformly distributed, so is the vector $\left(\left[\frac{\overline{b_i}}{\overline{f}}\right]_{-1}, \dots, \left[\frac{\overline{b_i}}{\overline{f}}\right]_{m-2}\right)$. Moreover, since the free coefficient of f is invertible in \mathbb{Z}_q , the matrix F is invertible modulo q and maps the uniform distribution to the uniform one and we get that, in particular, $\left(\left[\frac{\overline{b_i}}{\overline{f}}\right]_{m-1}, \dots, \left[\frac{\overline{b_i}}{\overline{f}}\right]_{m-1+d-1}\right)$ is also uniform. As a consequence, $b'_i(x)$ will also be uniform.

Notice that the only condition on f_m used in the reduction is the invertibility in \mathbb{Z}_q . Still, we keep the stronger condition $f_m = 1$ in the statement of the theorem just to be consistent with the definition of the $\mathsf{PLWE}^{(f)}$ problem.

4.7 Hardness of MP-LWE with small secrets

In the next chapter, we will show the cryptographic relevance of MP-LWE by builduing a public-key encryption scheme and a signature scheme whose proofs of security rely on the presumed hardness of MP-LWE. A main obstacle towards building a signature scheme directly from MP-LWE with the Fiat-Shamir with aborts methodology [Lyu09] is the need of small secrets. In this section we show that MP-LWE remains at least as hard as PLWE for numerous parametrizing polynomials f, when the secret s is sampled from a specific distribution χ_s that produces small secrets with overwhelming probability. We use the so-called "discrete variant" of MP-LWE, where the error distribution is over $\mathbb{Z}_q^{<d}[x]$ instead of $\mathbb{R}_q^{<d}[x]$. For the ease of reading, we keep the same notations for the MP-LWE problem with discrete noise as for the MP-LWE problem with continuous noise.

By $J_i \in \mathbb{Z}^{i \times i}$ we denote the matrix with 1's on the anti-diagonal and 0's everywhere else. Let $q \ge 2$, $n \ge d > 0, T > 0$ and k := n + d - 1. Let $\mathcal{E}(T, d, n)$ denote the set of all monic polynomials $g(x) \in \mathbb{Z}[x]$ with constant coefficient coprime to q, degree $m \in [d, n]$, and $\sigma_m(\mathbf{M}_f) \ge T$.

Theorem 4.3. Let $q \ge 2$, $n \ge d > 0$, T > 0 and k := n + d - 1. For any polynomial $f \in \mathcal{E}(T, d, n)$ and $1 \ge \alpha \ge \frac{2\sqrt{n}}{qT}$, there is a ppt reduction from $\mathsf{PLWE}_{q, D_{\mathbb{Z}^m, \alpha q}, D_{\mathbb{Z}^m, \alpha q}}^{(f)}$ to $\mathsf{MP-LWE}_{q, n, d, D_{\mathbb{Z}^d, \alpha''q}, D_{\mathbb{Z}^k, \alpha'q}}$, where $\alpha' = \alpha n \sqrt{2n} \cdot \mathsf{EF}(f)^2$ and $\alpha'' = \alpha \sqrt{2d} \cdot \mathsf{EF}(f)$.

Proof. We first reduce $\mathsf{PLWE}^{(f)}$ to a variant of $\mathsf{MP}\text{-}\mathsf{LWE}$ where the dependence on f lies both in the secret and error distributions. Using the same idea as in Theorem 4.1, except for the fact that now we do not rerandomize the secret to make it uniform, we know that there is a **ppt** reduction from $\mathsf{PLWE}_{q,\chi_e,\chi_s}^{(f)}$ to $\mathsf{MP}\text{-}\mathsf{LWE}_{q,n,d,\chi'_e,\chi'_s}$ where $\chi'_e = \mathsf{J}_d \cdot \mathsf{M}_f^d \cdot \chi_e$ and $\chi'_s = \mathsf{J}_{n+d-1} \cdot \mathsf{Rot}_f^{d+n-1}(1) \cdot \mathsf{M}_f \cdot \chi_s$. We now define the following notations: $B_s := \alpha \cdot q \mathsf{J}_k \cdot \mathsf{Rot}_f^k(1) \cdot \mathsf{M}_f$ and $B_e := \alpha \cdot q \mathsf{J}_d \cdot \mathsf{M}_f^d$, and $\Sigma_s := B_s \cdot B_s^t \in \mathbb{R}^{k \times k}$ and $\Sigma_e := B_e \cdot B_e^t \in \mathbb{R}^{d \times d}$, respectively. This means that there is a a **ppt** reduction from $\mathsf{PLWE}_{q,D_{\mathbb{Z}^m,\alpha q},D_{\mathbb{Z}^m,\alpha q}}^{(f)}$ to $\mathsf{MP}\text{-}\mathsf{LWE}_{q,n,d,D_{\mathbb{Z}^d,\Sigma_s}}, D_{\mathbb{Z}^k,\Sigma_s}$. We now have, using Lemmas 2.4 and 2.5, that

$$\begin{split} \|\Sigma_s\| &\leq (\alpha q)^2 \cdot \|\mathsf{Rot}_f^{d+n-1}(1)\|^2 \cdot \|\mathbf{M}_f\|^2 \\ &\leq (\alpha q)^2 \cdot \left(m + (d+n-1-m) \cdot m \cdot \mathsf{EF}(f)^2\right) m \cdot \mathsf{EF}(f)^2 \\ &\leq (\alpha q)^2 \cdot (n + (n-1) \cdot n \cdot \mathsf{EF}(f)^2) n \cdot \mathsf{EF}(f)^2 \\ &\leq (\alpha q)^2 \cdot n^3 \cdot \mathsf{EF}(f)^4 < (\alpha' q)^2/2 \end{split}$$

and

$$\|\Sigma_e\| \le (\alpha q)^2 \cdot \|\mathbf{M}_f^d\|^2 \le d \cdot (\alpha q \cdot \mathsf{EF}(f))^2 < (\alpha'' q)^2/2.$$

Since $\|\Sigma_s\| < (\alpha'q)^2$ and $\|\Sigma_e\| < (\alpha''q)^2$, there exist two symmetric positive definite matrices Σ'_s and Σ'_e such that $\Sigma_s + \Sigma'_s = (\alpha'q)^2 \mathbf{I}_k$ and $\Sigma_e + \Sigma'_e = (\alpha''q)^2 \mathbf{I}_d$. We now replace the rerandomization to uniform from Theorem 4.1 by a rerandomization to a Gaussian distribution. We first sample $t \leftrightarrow D_{\mathbb{Z}^k, \Sigma'_s}$. For any MP-LWE_{q,n,d,D_{Z^d, \Sigma_e}, D_{Z^k, \Sigma_s} sample (a_i, b_i) , we sample $e' \leftrightarrow D_{\mathbb{Z}^d, \Sigma'_e}$ and output $(a'_i, b'_i) = (a_i, b_i + a_i \odot_d t + e'_i)$. If (a_i, b_i) is uniform, so is (a'_i, b'_i) . If $b_i = a_i \odot_d s + e_i$, then}

$$b'_{i} = a_{i} \odot_{d} s + e_{i} + a_{i} \odot_{d} t + e'_{i} = a_{i} \odot_{d} (s + t) + (e_{i} + e'_{i}).$$

The matrices Σ_s , Σ'_s , Σ_e and Σ'_e are all symmetric, so they are in particular orthogonally diagonalizable. Moreover, since Σ_s and Σ'_s (resp. Σ_e and Σ'_e) commute, it means that Σ_s and Σ'_s (resp. Σ_e and Σ'_e) are simultaneously diagonalizable. We can hence write $\Sigma_s = UD_sU^t$ and $\Sigma'_s = UD'_sU^t$ for two diagonal matrices D_s and D'_s such that $(\alpha' q)^2 \mathbf{I}_k = D_s + D'_s$ and an orthogonal matrix $U \in \mathbb{R}^{k \times k}$. Similarly, we can write $\Sigma_e = VD_eV^t$ and $\Sigma'_e = VD'_eV^t$, where D_e and D'_e are diagonal, $D_e + D'_e = (\alpha''q)^2 \mathbf{I}_d$ and $V \in \mathbb{R}^{d \times d}$ is orthogonal. Now we can write

$$\eta_{2^{-k}}(\sqrt{\Sigma_s^{-1} + \Sigma_s'^{-1}} \cdot \mathbb{Z}^k) = \eta_{2^{-k}}(\sqrt{U(D_s^{-1} + D_s'^{-1})U^t} \cdot \mathbb{Z}^k) \qquad = \eta_{2^{-k}}(U\sqrt{D_s^{-1} + D_s'^{-1}} \cdot \mathbb{Z}^k).$$

Since the smoothing parameter is invariant to rotations, we have that

$$\eta_{2^{-k}}(\sqrt{\Sigma_s^{-1} + \Sigma_s'^{-1}} \cdot \mathbb{Z}^k) = \eta_{2^{-k}}(\sqrt{D_s^{-1} + D_s'^{-1}} \cdot \mathbb{Z}^k).$$

Using Lemma 2.9, we obtain that

$$\eta_{2^{-k}}(\sqrt{D_s^{-1} + D_s'^{-1}} \cdot \mathbb{Z}^k) \le \max_i \sqrt{\frac{1}{\sigma_i(\Sigma_s)} + \frac{1}{(\alpha' q)^2 - \sigma_i(\Sigma_s)}} \cdot \sqrt{k+1}.$$

We showed that $\sigma_1(\Sigma_s) \leq (\alpha q)^2 \sigma_1(\mathbf{M}_f)^2 \sigma_1(\mathsf{Rot}_f^{d+n-1}(1))^2 \leq (\alpha' q)^2/2$, which implies that $(\alpha' q)^2 - (\alpha' q)^2/2$

 $\sigma_i(\Sigma_s) \ge \sigma_i(\Sigma_s)$ for any $i \le k$ and thus for any $i \le k$,

$$\frac{1}{\sigma_i(\Sigma_s)} + \frac{1}{(\alpha' q)^2 - \sigma_i(\Sigma_s)} \le \frac{2}{\sigma_i(\Sigma_s)} \le \frac{2}{\sigma_k(\Sigma_s)}$$

Using the bound on the smallest singular value of \mathbf{M}_{f} , we now get that

$$\sigma_k(\Sigma_s) \ge (\alpha q)^2 \sigma_m(\mathbf{M}_f)^2 \sigma_m(\mathsf{Rot}_f^{n+d-1}(1))^2 \ge (\alpha q)^2 \cdot T^2,$$

which guarantees that for any $\alpha \geq \frac{2\sqrt{n}}{q \cdot T}$ we have that

$$\eta_{2^{-k}}(\sqrt{D_s^{-1} + D_s'^{-1}} \cdot \mathbb{Z}^k) \le \sqrt{\frac{2}{(\alpha q)^2 \cdot T^2}} \cdot \sqrt{k+1} \le 1.$$

As a consequence, using Lemma 2.7, the statistical distance between the distribution of s+t and $D_{\mathbb{Z}^k,\alpha' q}$ is $< 4 \cdot 2^{-d} = 4\varepsilon$ as k > d. Similarly, we have $\eta_{2^{-d}}(\sqrt{\Sigma_e^{-1} + \Sigma_e'^{-1}} \cdot \mathbb{Z}^d) \le 1$ and the statistical distance between the distribution of $e_i + e'_i$ and $D_{\mathbb{Z}^d,\alpha''q}$ is also $\le 4\epsilon$. This completes the proof.

Lemma 4.5. Let $f = x^m + P(x) \in \mathbb{Z}[x]$ with $m \ge 2$ and $\deg(P) \le m/2$. Then $\sigma_m(\mathbf{M}_f) \ge \frac{1}{2+\sqrt{m}\cdot\mathsf{EF}(f)}$. *Proof.* By reordering the rows of \mathbf{M}_f , the singular values stay the same and we can view \mathbf{M}_f as a block of four matrices $D_1 \in \mathbb{Z}^{\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor}$, $D_2 \in \mathbb{Z}^{\lceil m/2 \rceil \times \lceil m/2 \rceil}$, $0 \in \mathbb{Z}^{\lceil m/2 \rceil \times \lfloor m/2 \rfloor}$ and $T \in \mathbb{Z}^{\lfloor m/2 \rfloor \times \lceil m/2 \rceil}$ in the following way:

$$\mathbf{M}_f = \begin{bmatrix} D_1 & T \\ 0 & D_2 \end{bmatrix}.$$

The matrices D_1 and D_2 are diagonal, 0 is the all-0 matrix and T is an upper triangular matrix. We now use the definition $\sigma_m(\mathbf{M}_f) = \min(\|\mathbf{M}_f \cdot y\|_2 : y \in \mathbb{R}^m, \|y\|_2 = 1)$. Let $y \in \mathbb{R}^m$ such that $\sigma_m(\mathbf{M}_f) = \|\mathbf{M}_f \cdot y\|_2$ and $\|y\|_2 = 1$. The vector y can be written as $y = (y_0^t | y_1^t)^t$, with $y_0 \in \mathbb{R}^{\lfloor m/2 \rfloor}$ and $y_1 \in \mathbb{R}^{\lceil m/2 \rceil}$. On the one hand, we have:

$$\begin{split} \|\mathbf{M}_{f} \cdot y\|_{2} &\geq \|D_{1} \cdot y_{0} + T \cdot y_{1}\|_{2} &\geq \|D_{1} \cdot y_{0}\|_{2} - \|T \cdot y_{1}\|_{2} \\ &\geq \|y_{0}\|_{2} - \|T\| \cdot \|y_{1}\|_{2} \\ &\geq \|y\|_{2} - \|y_{1}\|_{2} - \|\mathbf{M}_{f}\| \cdot \|y_{1}\|_{2} \\ &\geq 1 - (1 + \sqrt{m} \cdot \mathsf{EF}(f)) \cdot \|y_{1}\|_{2}, \end{split}$$

where the last inequality is by Lemma 2.4. On the other hand, we also have

$$\|\mathbf{M}_f \cdot y\|_2 \ge \|D_2 \cdot y_1\|_2 \ge \|y_1\|_2.$$

This provides the bound

$$\sigma_m(\mathbf{M}_f) \ge \max\left(1 - (1 + \sqrt{m} \cdot \mathsf{EF}(f)) \cdot \|y_1\|_2, \|y_1\|_2\right) \ge \frac{1}{2 + \sqrt{m} \cdot \mathsf{EF}(f)},$$

and the conclusion follows.

An elementary computation shows that for any polynomial as in the above Lemma 4.5, we have $\mathsf{EF}(f) \leq \frac{3}{4}m^2 \|P\|_{\infty}^2$ (see also [LM06, Se. 3.1] for a similar but more general statement). This implies the following corollary of Theorem 4.3.

Corollary 4.1. Let $q \ge 2$, $n \ge d > 0$, k := n + d - 1 and S > 0. For any degree $m \ge 2$ polynomial $f = x^m + P(x) \in \mathbb{Z}[x]$ with constant coefficient coprime with q such that $\deg(P) \le m/2$ and $\|P\|_{\infty}^2 \le m/2$.

$$\begin{split} & 4S/3m^2 \ \text{and any } 1 \geq \alpha \geq 2\sqrt{n} \cdot (2+\sqrt{n}S)/q \ \text{there is a ppt reduction from } \mathsf{PLWE}_{q,D_{\mathbb{Z}^m,\alpha q},D_{\mathbb{Z}^m,\alpha q}}^{(f)} \ \text{to} \\ & \mathsf{MP-LWE}_{q,n,d,D_{\mathbb{Z}^d,\alpha''q},D_{\mathbb{Z}^k,\alpha' q}}, \ \text{where } \alpha' = \alpha n\sqrt{2n} \cdot S^2 \ \text{and } \alpha'' = \alpha \sqrt{2d} \cdot S. \end{split}$$

Chapter 5

Applications of MP-LWE in Cryptography

In this chapter, we exhibit the cryptographic expressiveness of MP-LWE by constructing a public-key encryption scheme and a digital signature scheme whose proofs of security rely on the hardness of PLWE for at least one polynomial f of degree n in a family whose size is exponentially large as a function of n. We also argue why the technique used in [GPV08] to build the so-called *dual* of the encryption scheme above cannot be applied in the MP-LWE setting.

This chapter is mainly based on two articles: a joint work [RSSS17] with Amin Sakzad, Damien Stehlé and Ron Steinfeld which was accepted at Crypto 2017 and a joint work [BDH+20] with Shi Bai, Dipayan Das, Ryo Hiromasa, Amin Sakzad, Damien Stehlé, Ron Steinfeld and Zhenfei Zhang accepted at PKC 2020.

Contents

5.1	Intr	oduction	77
5.2	Con	$\operatorname{tributions}$	77
	5.2.1	Related works	78
	5.2.2	Follow-up works	79
5.3	Ар	ublic-key encryption scheme from MP-LWE	80
	5.3.1	The scheme	80
	5.3.2	Correctness	80
	5.3.3	Security	81
	5.3.4	Parameters	83
5.4	An i	impossibility result for a dual-Regev scheme based on $MP-LWE\ldots$	83
5.5	A si	gnature scheme based on small secrets MP-LWE	84
	5.5.1	The identification scheme	84
	5.5.2	The signature scheme	88
	5.5.3	Concrete parameters	89
	5.5.4	Implementation	91
	5.5.5	An attack on Inhomogeneous $PSIS^\emptyset$ with small secrets $\hdots \hdots \hdot$	91

5.1 Introduction

The first cryptographic scheme based on LWE was the public-key encryption scheme given by Regev in [Reg09]. Assume n is the dimension and q is the modulus of the underlying LWE problem. The public key consists of m LWE samples $(\mathbf{a}_i, b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i \mod q) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$, while the secret key consists of the commun secret s used to generate the respective samples. To encrypt a bit $\mu \in \{0, 1\}$, one takes a random subset S of [m] and outputs the ciphertext $\mathbf{c} = (\sum_{i \in S} \mathbf{a}_i, \mu \cdot \lfloor q/2 \rfloor + \sum_{i \in S} b_i)$, which can be interpreted as an element in \mathbb{Z}_q^{n+1} . To decrypt **c**, one computes $(\mathbf{s}, 1)^t \cdot \mathbf{c}$ and decrypts to 0 if this is closer to 0 than to |q/2| modulo q or decrypts to 1 otherwise. If the errors are small compared to q, than the decryption is correct with high probability. The security proof is based on two observations: the public key is indistinguishable from uniform by the LWE assumption and encrypting using a uniform key is information-theoretically secure. A major drawback of the above cryptosystem is that it allows the encryption of only one bit at a time. In order to fix this, Peikert et al. described in [PVW08] an improved version of the above cryptosystem, which allows the encryption of l bits at a time. The main idea is to replace the secret $\mathbf{s} \in \mathbb{Z}_q^n$ with a matrix $\mathsf{S} \in \mathbb{Z}_q^{l \times n}$ whose rows represent l uniformly independent LWE secrets. The encryption scheme allows to encrypt l = O(n) bits per ciphertext, with no asymptotic increase in the sizes of the public key or ciphertexts, nor in the runtime of encryption at the cost of an increase in the size of the secret key size and decryption runtime.

Later on, Gentry, Peikert and Vaikuntanathan described in [GPV08] a public-key encryption scheme which can be seen as a "dual" of the initial Regev cryptosystem, in the sense that the key generation and the encryption steps are swapped. In the dual scheme, the secret key is an uniformly vector $\mathbf{x} \in \{0, 1\}^m$ and the public key consists of m uniform vectors $\mathbf{a}_i \in \mathbb{Z}_q^n$ and $\mathbf{u} = \sum_{i \in [m]} \mathbf{a}_i x_i$. To encrypt a bit $\mu \in \{0, 1\}$, one chooses a secret $\mathbf{s} \in \mathbb{Z}_q^n$ and outputs the ciphertext $\mathbf{c} = (\mathbf{s}^t \cdot \mathbf{a}_1, \dots, \mathbf{s}^t \cdot \mathbf{a}_m, \mathbf{s}^t \cdot \mathbf{u} + \mu \cdot \lfloor q/2 \rfloor)$. To decrypt, one computes $\mathbf{c}^t \cdot (\mathbf{x}, 1)$ and decrypts to 0 if this value is closer to 0 than to $\lfloor q/2 \rfloor$ and to 1 otherwise. The security proof has two main ideas: for a reasonable choice of parameters, the public key is nearly uniform using the leftover hash lemma and a public key along with a ciphertext is indistinguishable from uniform by the LWE hardness assumption.

The above LWE based encryption schemes also enjoy PLWE and RLWE counterparts ([SSTX09, LPR10, LPR13], among others).

The SIS problem and its variants allow the construction of digital signatures secure both in the standard model ([CHKP10, Boy10, MP12, DM14], etc.) and random-oracle model ([GPV08, Lyu09, Lyu12, GLP12, DDLL13], etc.).

5.2 Contributions

In this chapter, we illustrate the cryptographic expressiveness of MP-LWE, but also its limitations.

In Section 5.3, we first describe a public-key encryption scheme adapted from [Reg09] to the middleproduct setting. The scheme is IND-CPA secure under the MP-LWE hardness assumption, involves keys of bit-size $\tilde{O}(\lambda)$ and algorithms running in time $\tilde{O}(\lambda)$. The correctness proof of the scheme uses the associativity property of the middle product. To establish its security, we prove that a related hash function family involving middle-products is universal and apply the generalized version of the leftover hash lemma from Section 2.3.2. The standard leftover hash lemma does not seem to suffice for our needs, as in our case the first part of the ciphertext is not statistically close to uniform as it happens in Regev's encryption scheme.

In Section 5.5, we build an identification scheme which follows Schnorr's general framework [Sch89] and then upgrade it to a signature scheme MPSign that is tightly secure in the quantum-access random oracle model, using [KLS18]. The signature scheme is secure under the MP-LWE with short secrets hardness assumption. We show that MPSign is UF-CMA, which means that no adversary may forge

a signature on a message for which it has not seen a signature before. We did not manage to prove that there is no adversary who may forge a new signature on a previously signed message, *i.e.*, that the scheme is UF-sCMA. Nevertheless, as discussed in Chapter 2, any UF-CMA secure signature can be upgraded to a UF-sCMA secure signature using a one-time UF-sCMA secure signature [Kat10, p. 27].

We provide concrete parameters for MPSign corresponding to level 1 security of the NIST postquantum standardization process (via the SVP core hardness methodology from [ADPS16]), which take into account our tight quantum random-oracle model security proof with respect to small secret MPLWE (rather than just taking in account the classical random-oracle model security proof as, e.g., in the Dilithium scheme parameter selection [DKL⁺18]). We also provide parameters that achieve similar security to those from [Lyu16], to allow for a reasonably fair comparison. The MPSign verification key is larger but its signature size is twice smaller.

Our MPSign signature length savings over the scheme of [Lyu16] arise mainly due to our use of much smaller secret key coordinates. Therefore, one could wonder the reducing the size of the secret key coordinates in the scheme of [Lyu16] would also give a secure signature scheme. As an additional small contribution we show that the answer is negative by presenting a simple efficient key recovery attack on Lyubashevsky's scheme with sufficiently small secret coordinates. Our attack works (heuristically) when the underlying inhomogeneous variant of $PSIS^{\emptyset}$ has a unique solution, and shows that a lower bound similar to that shown sufficient in the security proof of [Lyu16] is also *necessary* for the security of Lyubashevsky's scheme (and the underlying inhomogeneous PSIS^{\emptyset} problem) with small secret coordinates. Zhenfei Zhang implemented a proof-of-concept of MPSign in Sage, and the code is publicly available at:

https://github.com/pqc-ntrust/middle-product-LWE-signature.

Adapting the dual-Regev scheme from [GPV08] does not seem straightforward. Indeed, it appears that we would need a leftover hash lemma for polynomials over $\mathbb{Z}_q[x]$ that are not folded modulo some polynomial f. The difficulty is that the constant coefficients of the polynomials are now "isolated", in the sense that the constant coefficient of a polynomial combination of polynomials only involves the constant coefficients of these polynomials. At the end of this chapter we provide evidence that even if we assume the existence of such a leftover hash lemma, the technique used in [GPV08] to prove correctness cannot be applied to a direct adaptation of the respective scheme to the middle-product setting.

5.2.1 Related works

Our signature construction is similar to the one in [Hir18]. However, the proof of the latter is incorrect: in its proof of high min-entropy of commitments (see [Hir18, Lemma 7]), it is assumed that the middle n coefficients of the product between a uniform $a \in \mathbb{Z}_q[x]$ of degree < n and a fixed polynomial y of degree $\leq 2n$, are uniform. In fact, this distribution depends on the rank of a Hankel matrix associated to y and encoding the linear function from a to the considered coefficients of the product. This Hankel matrix can be of low rank and, when it is the case, the resulting distribution is uniform on a very small subset of the range. Interestingly, the distribution of these Hankel matrices (for a uniform y) was recently studied in [BBD+19], in the context of proving hardness of an MP-LWE variant with deterministic noise. We do not know how to fix the error from [Hir18]. As a result, we use a different identification scheme to be able to make our proofs go through. Concretely, the identification scheme from [Hir18] used the Bai-Galbraith [BG14] compression technique to decrease the signature size. We circumvent the difficulty by not using the Bai-Galbraith compression technique.

Lyubashevsky's signature from [Lyu16] can also be viewed as secure under the assumption that $\mathsf{PLWE}^{(f)}$ is hard for at least one f among exponentially many defining polynomials f, like ours. Indeed, it was proved secure under the assumption that $\mathsf{PSIS}^{\emptyset}$ is hard, it was proved that $\mathsf{PSIS}^{(f)}$ reduces

to $\mathsf{PSIS}^{\emptyset}$ for exponentially many defining polynomials f, and $\mathsf{PLWE}^{(f)}$ (directly) reduces to $\mathsf{PSIS}^{(f)}$. Furthermore, $\mathsf{MP-LWE}$ (both with small-magnitude secrets and uniform secrets) reduces to $\mathsf{PSIS}^{\emptyset}$, whereas the converse is unknown. Hence it seems that in terms of assumptions, Lyubashevsky's signature outperforms ours. However, the security proof from [Lyu16] only holds in the random oracle model, as opposed to ours which is tight in the quantum-access random oracle model. Recent techniques on Fiat-Shamir in the QROM [LZ19, DFMS19] might be applicable to [Lyu16], but they are not tight.

It is useful to also compare MPSign with LWE-based signature schemes and efficient lattice-based signature schemes such as those at Round 2 of the NIST post-quantum standardization process [NIS]: Dilithium [DKL⁺18], Falcon [PFH⁺19] and Tesla [BAA⁺19]. Compared to LWE-based signatures, our proposal results in much smaller values for the sum of sizes of a signature and a public key, with much stronger security guarantees than the efficient schemes based on polynomial rings. For example, scaling Dilithium with NIST security level 1 parameters to LWE requires multiplying the public key size by the challenge dimension n = 256, since for an LWE adaptation of Dilithium, the public key would be a matrix with n columns instead of 1. For NIST security level 1, the public key and signature sizes sum would be above 300KB for an LWE adaptation of Dilithium, whereas the same quantity is 47KB for MPSign (see Table 5.2). Now, compared to the Dilithium, Falcon and Tesla NIST candidates, security guarantees are different. The security of Dilithium and Tesla relies on the module variants of PLWE and PSIS for a fixed polynomial [LS15]. In the case of Dilithium, the known security proof in the QROM is quite loose [LZ19], unless one relies on an ad hoc assumption like SelfTargetMSIS [KLS18]. Moreover, in the case of Dilithium, the SIS instance is in an extreme regime: the maximum infinity norm of the vectors to be found are below q/2, but their Euclidean norms may be above q. Currently, no reduction backs the assumption that SIS is intractable in that parameter regime. In Falcon, the public key is assumed pseudo-random, which is an adhoc version of the NTRU hardness assumption [HPS98]. Oppositely, the security of MPSign relies on the assumed PLWE hardness for at least one polynomial among exponentially many. Overall, MPSign is an intermediate risk-performance tradeoff between fixed-ring and LWE-based schemes.

5.2.2 Follow-up works

In this section, we briefly describe two works built upon MP-LWE.

5.2.2.1 Titanium: KEM from the MP-LWE assumption

Titanium ([SSZ17, SSZ19]) is an optimized public-key encryption scheme built upon our results from Section 5.3 and submitted to the NIST standardization competition [NIS]. In Titanium, Steinfeld *et al.* first specialize the hardness result on MP-LWE to a restricted (but still exponentially large) family of polynomials \mathcal{F} . For any $f \in \mathcal{F}$, the PLWE^(f) to MP-LWE reduction preserves the noise distribution.

The proposal specifies two variants of Titanium: Titanium-CPA and Titanium-CCA. Titanium-CPA is tightly IND-CPA secure in the random-oracle model under the $\mathsf{PLWE}^{(f)}$ hardness assumption with respect to any f in the family \mathcal{F} . Titanium-CCA is obtained by converting Titanium-CPA into an IND-CCA secure Key Encapsulation Mechanism (KEM) using the generic Fujisaki-Okamoto transformation [FO99, HHK17]. In the classical setting, the Fujisaki-Okamoto conversion is tight and the concrete parameters of Titanium are chosen by taking into account proof bounds and the best known BKZ attack on the underlying $\mathsf{PLWE}^{(f)}$ problem associated to a polynomial $f \in \mathcal{F}$ of maximum degree and following the CoreSVP methodology from [ADPS16]. In the quantum setting, the parameters of Titanium are chosen based on the assumption that the classical security proof bounds still apply.

The Std128 parameter set of Titanium-CPA corresponds to NIST category 1 or AES 128 security level. At a higher security level, Titanium-CPA instantiated with Std128 has ciphertexts 3 times smaller, faster key generation, encryption, and decryption time by factors of 1.4, 2.3, and 1.3 and

shorter |pk| + |ct| size compared to the LWE based encryption scheme described in [BCD⁺16]. On the other hand, compared to the Module-LWE based scheme Kyber [BDK⁺18], at a lower security level, Titanium-CCA has ciphertexts, secret key, and public key that are at least 3 times larger and key generation, encapsulation, and decapsulation times slower by factors of 7.6, 5 and 5.1. On the security front, Titanium achieves better security guarantees than Kyber, by not relying on the choice of a specific polynomial. Consequently, Titanium could be seen as an intermediate solution in terms of security guarantees versus efficiency.

5.2.2.2 Identity-based encryption from MP-LWE

Lombardi *et al.* [LVV19] introduced a slight variant of MP-LWE, the Degree-Parametrized-MP-LWE problem, which remains at least as hard as the PLWE problem for a large class of polynomials f. In this new variant, the samples generated using a fixed secret polynomial s can have varying pre-specified degrees. They also proved a leftover-hash lemma for polynomials with bounded degree and proposed a so-called *dual-Regev type* encryption scheme based on the Degree-Parametrized-MP-LWE assumption. This is a variant of the Regev type encryption we build in this chapter, where the key generation and ecryption steps are swapped. Their scheme is IND-CPA secure in the random oracle model under the Degree-Parametrized-MP-LWE assumption and has quasi-linear key size and algorithm runtime.

Their main contribution is an Identity-Based Encryption (IBE) scheme based on Degree-Parametrized-MP-LWE. Their construction follows the lattice trapdoors paradigm of [GPV08] and is obtained by combining the dual encryption scheme with Micciancio-Peikert style lattice trapdoors [MP12]. The IBE scheme is (T, ε) secure under the (T, ε) Degree-Parametrized-MP-LWE assumption only for $\varepsilon > 2^{-\text{poly}(\log n)}$. This technical limitation is due to the achievable parameters of their leftover hash lemma. Since for exponential security, one needs to be able to handle exponentially small ε , their IBE scheme does not reach a meaningful form of concrete security.

5.3 A public-key encryption scheme from MP-LWE

In this section we describe a public key encryption scheme that is IND-CPA secure under the MP-LWE hardness assumption. The scheme is an adaptation of Regev's from [Reg09] and can be upgraded to an IND-CCA secure scheme in the random-oracle model using the Fujisaki-Okamoto [FO99] transform.

5.3.1 The scheme

The public-key encryption scheme we propose in Figure 5.1 relies on parameters $q, n, d, t \geq 2$ with q odd, and a noise rate $\alpha \in (0, 1)$. We let $\chi = \lfloor D_{\alpha q} \rceil$ denote the distribution over $\mathbb{Z}^{\leq d+k}[x]$ where each coefficient is sampled on \mathbb{R} from $D_{\alpha \cdot q}$ and then rounded to nearest integer. The plaintext space is $\{0, 1\}^{\leq d}[x]$ and the ciphertext space is $\mathbb{Z}_q^{\leq k+n}[x] \times \mathbb{Z}_q^{\leq d}[x]$.

5.3.2 Correctness

The correctness of the scheme presented in Figure 5.1 follows from Lemma 4.2 and the proof of correctness of Regev's encryption scheme.

Lemma 5.1. Assume that $\alpha < 1/(16\sqrt{\lambda tk})$ and $q \ge 16t(k+1)$. With probability $\ge 1 - d \cdot 2^{-\Omega(\lambda)}$ over the randomness of $(sk, pk) \leftarrow KeyGen$, for all plaintext m and with probability 1 over the randomness of Enc, we have Dec(sk, Enc(pk, m)) = m.

 $KevGen(1^{\lambda})$ Enc(pk, m)1: $s \leftarrow U(\mathbb{Z}_q^{< n+d+k-1}[x])$ 1: i = 12: for $i \leq t$ do 2: i = 13: $r_i \leftrightarrow U(\{0,1\}^{\leq k+1}[x])$ 3: for i < t do 4: 5: 6: 7: end for 8: sk = s9: $\mathsf{pk} = (a_i, b_i)_{i \le t}$ Dec(sk, c)10: output (pk, sk) 1: $m' = (c_2 - c_1 \odot_d s \mod q \mod 2$ 2: output m'

Fig. 5.1: PKE scheme from MP-LWE.

Proof. Assume that (c_1, c_2) is an encryption of m under pk. Then we have, modulo q:

$$c_{2} - c_{1} \odot_{d} s = m + \sum_{i \leq t} r_{i} \odot_{d} b_{i} - \left(\sum_{i \leq t} r_{i} \cdot a_{i}\right) \odot_{d} s$$

$$= m + \sum_{i \leq t} \left(r_{i} \odot_{d} \left(a_{i} \odot_{d+k} s + 2 \cdot e_{i}\right) - \left(r_{i} \cdot a_{i}\right) \odot_{d} s\right)$$

$$= m + 2\sum_{i \leq t} r_{i} \odot_{d} e_{i},$$

where the last equality follows from Lemma 4.2. If $||m + 2 \cdot \sum_{i \leq t} r_i \odot_d e_i||_{\infty} < q/2$, then centered reduction modulo q of $c_2 - c_1 \odot_d s$ gives us $m + 2 \cdot \sum_{i \leq t} r_i \odot_d e_i$ (over the integers). Reducing modulo 2 then provides m.

Now, each coefficient of $\sum_{i \leq t} r_i \odot_d e_i$ can be viewed as an inner product between a binary vector of dimension t(k+1) and a vector sampled from $\lfloor D_{\alpha q} \rceil^{t(k+1)}$. Each coefficient individually has magnitude $\leq \alpha q \sqrt{\lambda t(k+1)} + t(k+1)$ with probability $\geq 1 - 2^{-\Omega(\lambda)}$, because of the Gaussian tail bound and the triangle inequality. By the union bound and triangular inequality, we obtain that $\|m+2 \cdot \sum_{i \leq t} r_i \odot_d e_i\|_{\infty} < 2\alpha q \sqrt{t\lambda(k+1)} + 2t(k+1) + 1$ with probability $\geq 1 - d \cdot 2^{-\Omega(\lambda)}$.

5.3.3 Security

The security proof of the scheme presented in Figure 5.1 is adapted from that of Regev's encryption scheme from [Reg09], with a subtlety in the application of the leftover hash lemma. In Regev's scheme, if the public key is replaced by uniformly random elements, then the leftover hash lemma guarantees that the joint distribution of the public key and the encryption of an arbitrary plaintext is within exponentially small statistical distance from uniform. This property does not hold in our case: indeed, if a_1, \ldots, a_t all have constant coefficient equal to 0 (this event occurs with a probability $1/q^t$, which is not exponentially small for our parameters), then so does $\sum_i r_i a_i$. However, we can show that the second component c_2 of the ciphertext is statistically close to uniform, given the view of the first component c_1 . This suffices, as the plaintext is embedded in the second ciphertext component.

We first prove that the hash function family coming into play in the security proof is universal.

Lemma 5.2. Let $q, k, d \geq 2$. For $(b_i)_i \in (\mathbb{Z}_q^{\leq d+k}[x])^t$, we let $h_{(b_i)_i}$ denote the map that sends $(r_i)_{i\leq t} \in (\{0,1\}^{\leq k+1}[x])^t$ to $\sum_{i\leq t} r_i \odot_d b_i \in \mathbb{Z}_q^{\leq d}[x]$. Then the hash function family $(h_{(b_i)_i})_{(b_i)_i}$ is universal.

Proof. Our aim is to show that for r_1, \ldots, r_t not all 0, we have

$$\Pr_{(b_i)_i, (b'_i)_i} \left[\sum_{i \le t} r_i \odot_d b_i = \sum_{i \le t} r_i \odot_d b'_i \right] = q^{-d}.$$

W.l.o.g. we may assume that $r_1 \neq 0$. By linearity, it suffices to prove that for all $y \in \mathbb{Z}_q^{\leq d}[x]$,

$$\Pr_{b_1}\left[r_1 \odot_d b_1 = y\right] = q^{-d}.$$

Let j be minimal such that the coefficient in x^j of r_1 is non-zero (*i.e.*, equal to 1 as r_1 is binary). Then the equation $r_1 \odot_d b_1 = y$ restricted to entries j + 1 to j + d is a triangular linear system in the coefficients of b_1 with diagonal coefficients equal to 1. The map $b_1 \mapsto r_1 \odot_d b_1$ restricted to these coefficients of b_1 is hence a bijection. This gives the equality above.

Lemma 5.3. Assume that $t \ge (2 \cdot \lambda + (k + d + n) \cdot \log q)/(k + 1)$. Then the scheme above is IND-CPA secure, under the MP-LWE_{q,n,d+k,D_{aq}} hardness assumption.

Proof. Recall that in the IND-CPA security experiment, the challenger C first gives pk to the adversary \mathcal{A} . Then, \mathcal{A} sends back two plaintexts $m_0 \neq m_1$. Now the challenger samples a bit $b \leftrightarrow \{0, 1\}$, computes $c \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b)$ and sends c to \mathcal{A} , who eventually outputs a bit b'. The scheme is secure if no ppt adversary \mathcal{A} outputs b' = b more probability that is non-negligibly away from 1/2.

Now, consider the variant of the experiment above, in which C does not run $(\mathsf{sk}, \mathsf{pk}) \leftrightarrow \mathsf{KeyGen}(1^{\lambda})$ but instead samples $\mathsf{pk} = (a_i, b_i)_i$ uniformly. Under the MP-LWE hardness assumption, the probabilities that \mathcal{A} outputs b' = b in both experiments are negligibly close. The reduction from MP-LWE to distinguishing the first and second experiments consists in multiplying by 2 (which is co-prime to q) and rounding the real samples given by an MP-LWE oracle to the nearest integer modulo q. The latter maps MP-LWE with real noise to MP-LWE with rounded real noise (and uniform MP-LWE over the reals modulo q to a uniform MP-LWE over the integers modulo q).

We consider a third experiment, in which C also samples $\mathsf{pk} = (a_i, b_i)_i$, and additionally does not compute $c \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b)$ before sending it to \mathcal{A} , but instead computes $c = (c_1, c_2)$ as follows. For $i \leq t$, it samples $r_i \leftarrow U(\{0, 1\}^{\leq k+1}[x]), u \leftarrow U(\mathbb{Z}_q^{\leq d}[x])$, and sets:

$$c_1 = \sum_{i \le t} r_i \cdot a_i , \quad c_2 = u.$$

Note that in this game, the view of \mathcal{A} is independent of b, and hence the probability that it outputs b' = b is exactly 1/2. We argue below that the distributions of $((a_i, b_i)_i, c_1, c_2)$ in this new experiment and the latter one are within exponentially small statistical distance. The combination of these two facts provides the result.

It remains to prove that

$$\Delta\Big(((a_i, b_i)_i, \sum_{i \le t} r_i \cdot a_i, \sum_{i \le t} r_i \odot_a b_i) , \quad ((a_i, b_i)_i, \sum_{i \le t} r_i \cdot a_i, u)\Big) \le 2^{-\lambda},$$

where the a_i 's, b_i 's, r_i 's and u are uniformly sampled in $\mathbb{Z}_q^{<n}[x]$, $\mathbb{Z}_q^{<d+k}[x]$, $U(\{0,1\}^{<k+1}[x])$ and $\mathbb{Z}_q^{<d}[x]$, respectively. By Lemma 5.2, the hash function family $h_{(b_i)_i}$ is universal. Further, the quantity $\sum_{i\leq t}r_i \cdot a_i$ belongs to $\mathbb{Z}_q^{<k+n}$, of cardinality q^{k+n} . Hence, by the Generalized Leftover Hash Lemma (see Lemma 2.6), the statistical distance above is bounded from above by $(2^{-(k+1)\cdot t} \cdot q^{k+d+n})^{1/2}/2$.

5.3.4 Parameters

Example parameters are $n \ge \lambda$, k = d = n/2, $q = \Theta(n^{5/2+c}\sqrt{\log n})$, $t = \Theta(\log n)$ and $\alpha = \Theta(1/n\sqrt{\log n})$, for c > 0 arbitrary. For these parameters, the scheme is correct (by Lemma 5.1) and secure under MP-LWE_{$q,n,n,D_{\alpha q}$} (by Lemma 5.3). These parameters allow to rely on the assumed hardness of $\mathsf{PLWE}_{q,D_{\beta,q}}^{(f)}$ via Theorem 4.1, for $\beta = \Omega(\sqrt{n}/q)$ (hence preventing attacks à la [AG11]) and for any f monic of degree n, with constant coefficient coprime with q and expansion factor $\le n^c$. Finally, note that the scheme encrypts and decrypts n plaintext bits in time $\widetilde{O}(n)$, and the key pair has bit-length $\widetilde{O}(n)$.

5.4 An impossibility result for a dual-Regev scheme based on MP-LWE

In this section we show that it is impossible to naturally adapt the dual-Regev scheme based on LWE from [GPV08] and its proofs of correctness and security to the middle-product setting if the distribution of the randomness used in the public key generation has enough entropy. The security proof of such a cryptosystem would follow two steps. Firstly, we would have to show that the public key is statistically close to uniform. Secondly, we would have to argue that the encryption of a message is indistinguishable from uniform under the MP-LWE hardness assumption.

We are interested in adapting the dual-Regev scheme to the middle-product setting because following [GPV08], combining the dual-Regev scheme with the so-called *lattice trapdoors* yields identity-based encryption. Identity-based encryption (IBE) is a type of public-key encryption in which users can generate their public keys from some public identifiers (such as their email addresses). The secret key of an user is generated by a trusted authority using public information and secret trapdoors. One of the main advantages of an IBE scheme is the elimination of the need to predistribute the public keys of the users.

We present in Figure 5.2 our attempt of dual-Regev public key encryption scheme based on MP-LWE. We argue now the design rationale and the choice of parameters. Assume that we want to prove that the encryption of a message is indistinguishable from uniform under the MP-LWE_{q,n+1,m+1,\chi} hardness assumption, where q is the modulus, $n, m \ge 0$ and χ is the error distribution on $\mathbb{Z}^{m+1}[x]$ used to generate the MP-LWE samples. As a consequence, the degree of a_i is less than n + 1, the degree of the secret s is less than n + m - 1, the plaintext space is $\{0,1\}^{\leq m+1}[x]$ and the ciphertext space is $(\mathbb{Z}_q^{\leq m+1}[x])^{t+1}$, where t+1 is the number of MP-LWE samples. We assume that the randomnesses r_i used in the key generation are sampled from a distribution R on $\mathbb{Z}^{\leq d+1}[x]$ for some d > 0. In order to create a polynomial u of degree n to be used in the encryption process, we consider that the natural way would be to first create the polynomial $\sum_i r_i a_i$ and then extract n consecutive coefficients out of it. As a consequence, we set $u = [\sum_i r_i a_i]_{l_0}^{l_0+n} \in \mathbb{Z}_q^{\leq n+1}[x]$, for some $l_0 \ge 0$. For the decryption, we consider that the natural way would be to extract from $\sum_i r_i c_i$ a number of m + 1 consecutive coefficients from the kth to the (k + m)th position for some $k \ge 0$.

Theorem 5.1. (Informal) If $H_{\infty}(R) > \log(3)$ and assuming that $[\sum_{i} r_{i}a_{i}]_{l_{0}}^{l_{0}+n}$ is uniform in $\mathbb{Z}_{q}^{< n+1}[x]$ conditioned on the a_{i} 's, no matter how we choose $l_{0}, k, d \geq 0$, the scheme from Figure 5.2 is not correct.

The proof of the theorem uses the following result.

Lemma 5.4 (Schwartz-Zippel, [GV12]). Let $F \in \mathbb{F}_q[X_1, \ldots, X_n]$ be a non-zero polynomial of degree d. For any $i \in \{1, \ldots, n\}$, let P_i be a probability distribution on \mathbb{Z}_q such that for any i, $H_{\infty}(P_i) \ge \log(q) - h$ for some common fixed $0 \le h \le \log(q)$. If $x_i \leftarrow P_i$, then $\Pr[F(x_1, \ldots, x_n) = 0) \le 2^h \cdot d/q$.

Proof of Theorem 5.1. Since $[\sum_{i} r_{i}a_{i}]_{l_{0}}^{l_{0}+n}$ is uniform in $\mathbb{Z}_{q}^{< n+1}[x]$, we should have $l_{0} \leq d$ (otherwise, the coefficients from d+n to $l_{0}+n$ will always be zero). We show that the decryption does not work with

```
\mathsf{KevGen}(1^{\lambda})
                                                                                            Enc(pk, \mu)
                                                                    1: s \leftarrow U(\mathbb{Z}_q^{< n+m+1})
1: i = 1
2: for i \leq t do
                                                                    2: i = 1
3: a_i \xleftarrow{} U(\mathbb{Z}_q^{< n+1}[x])
                                                                    3: for i < t + 1 do
4:
    r_i \leftrightarrow R
                                                                           e_i \leftarrow \chi
                                                                    4:
5: end for
                                                                    5: end for
6: u = [\sum_{i} r_i a_i]_{l_0}^{l_0+n} \in \mathbb{Z}_q^{< n+1}[x]
                                                                    6: i = 1
7: sk := (r_i)_{i \le t}
                                                                    7: for i \leq t do
                                                                           c_i = a_i \odot_{m+1} s + 2e_i
8: \mathsf{pk} := ((a_i)_{i < t}, u)
                                                                    8:
9: output (pk, sk)
                                                                    9: end for
                                                                  10: c_{t+1} = u \odot_{m+1} s + 2e_{t+1} + \mu
                                                                  11: output c = (c_1, \ldots, c_t, c_{t+1})
                                                                                            \mathsf{Dec}(\mathsf{sk}, c)
```

1: $\mu' = (c_{t+1} - [\sum_{i \leq t} r_i \cdot c_i]_k^{k+m} \mod q) \mod 2$ 2: output μ'

Fig. 5.2: Attempt of a dual-Regev PKE scheme from MP-LWE.

high probability over the choices of r, a and s as it works in the dual-Regev scheme. In order to achieve correctness, the equality $[[ra]_{l_0}^{l_0+n} \cdot s]_n^{n+m} = [r \cdot [as]_n^{n+m}]_k^{k+m}$ should hold with high probability over the choices of r, a and s. Let us denote by $P_{r,a,s}(x)$ the polynomial $[[ra]_{l_0}^{l_0+n} \cdot s]_n^{n+m} - [r \cdot [as]_n^{n+m}]_k^{k+m}$ of degree m. Now notice that $\Pr_{r,a,s}[P_{r,a,s}(x) = 0] \leq \Pr_{r,a,s}[[P_{r,a,s}]_0 = 0]$ and the coefficients of the polynomial $P_{r,a,s}(x)$ are degree 3 polynomials in the coefficients of the r, a and s. Suppose that $[P_{r,a,s}]_0 = [[ra]_{l_0}^{l_0+n} \cdot s]_n - [r \cdot [as]_n^{n+m}]_k$, which is a polynomial of degree 3 in the coefficients of the polynomials r, a and s, is the zero polynomial. By plugging in $r = x^{l_0}$, $a = x^n$ and s = 1, we get that $k = l_0$. By taking $r = x^{l_0-1}$, a = 1 and $s = x^{n+1}$, we get that the polynomial evaluates to -1, which is a contradiction. In conclusion, we obtain that $[P_{r,a,s}]_0 = 0] \leq 2^h \cdot 3/q < 1$ for any $h < \log(q/3)$. \Box

5.5 A signature scheme based on small secrets MP-LWE

In this section, we build an identification scheme based on the middle-product learning with errors with small secrets assumption. Then, we show that Theorem 2.3 is applicable to our construction by checking all the theorem assumptions, as in [KLS18]. As a consequence, by the Fiat-Shamir transformation, we obtain a digital signature scheme that is secure under the middle-product learning with errors with small secrets assumption in the quantum random-oracle model.

5.5.1 The identification scheme

We first present in Figure 5.3 an identification scheme which makes use of the middle-product of polynomials.

We use an extendable output function Sam, *i.e.*, a function on bit strings in which the output can be extended to any required length. If we want the deterministic output y of Sam on input x to be uniformly distributed on the set S, we write $y \leftrightarrow S := Sam(x)$.

The key generation starts by choosing a random string ρ and expanding it into a uniform polynomial $a \in \mathbb{Z}_q^{<n}[x]$ using the function *Sam*. The public key consists of a sample (a, b) drawn from the $\mathsf{MP}_{q,n,d+k,\chi}(s)$ distribution, where both the secret s and the error e follow a Gaussian distribution of parameter $\alpha' q$, respectively $\alpha'' q$.

In the first step of the protocol, the prover chooses two polynomials y_1 and y_2 whose coefficients are bounded in absolute value by a', respectively a'', and sends to the verifier the polynomial $w = a \odot_d y_1 + y_2$. The verifier chooses a random challenge from the challenge space

$$D_H := \{ c \in \{0, 1, -1\}^{< k+1} [x] \text{ with } \|c\|_1 = \kappa \}$$

and sends it back to the prover. The challenge space consists of polynomials of small norms and the parameter κ is chosen such that the cardinality of the challenge space is large. The prover now applies rejection in order to make sure that his answer doesn't leak information about the secret key. Concretely, the prover computes $z_1 = c \odot_{n+d-1} s + y_1$ and $z_2 = c \odot_d e + y_1$ and checks if $||z_1||_{\infty} \leq A'$ and $||z_2||_{\infty} \leq A''$. If so, it accepts to send his answer (z_1, z_2) to the verifier. Otherwise, it aborts. We provide concrete parameters with which our scheme can be instantiated in practice in the next section.

Fig. 5.3: The identification scheme (IGen, $V, P = (P_1, P_2)$).

Lemma 5.5. If $A' + \|c \odot_{n+d-1} s\|_{\infty} \leq a'$ and $A'' + \|c \odot_d e\|_{\infty} \leq a''$, then the identification scheme is perfectly na-HVZK, i.e., its transcripts are publicly simulatable and $\varepsilon_{zk} = 0$.

Proof. Figure 5.4 (left) shows how to generate a real transcript using the secret key sk, and Figure 5.4 (right) shows how to simulate a transcript using only the public key pk. The identification scheme is perfectly na-HVZK if every pair of polynomials $(z_1, z_2) \in \mathbb{Z}_{\leq A'}^{\leq n+d-1}[x] \times \mathbb{Z}_{\leq A''}^{\leq d}[x]$ has the same probability to be generated in the Trans algorithm as in the Sim algorithm. This is indeed the case: our choice of parameters guarantees that $z_1 - c \odot_{n+d-1} s \in \mathbb{Z}_{\leq a'}^{\leq n+d-1}[x]$ and $z_2 - c \odot_d e \in \mathbb{Z}_{\leq a''}^{\leq d}[x]$ and moreover, for any secret key (s, e) and any pair (z_1, z_2) , we have that

$$\Pr(z_1 = c \odot_{n+d-1} s + y_1 | y_1 \leftrightarrow \mathbb{Z}_{\leq a'}^{< n+d-1}[x]) = \Pr(y_1 = z_1 - c \odot_{n+d-1} s | y_1 \leftrightarrow \mathbb{Z}_{\leq a'}^{< n+d-1}[x])$$

and

$$\Pr(z_2 = c \odot_d e + y_2 | y_2 \leftrightarrow \mathbb{Z}_{\leq a''}^{\leq d}[x]) = \Pr(y_2 = z_2 - c \odot_d s | y_2 \leftrightarrow \mathbb{Z}_{\leq a''}^{\leq d}[x]).$$

As a consequence, the probability of producing z_1 and z_2 in Trans such that $||z_1||_{\infty} \leq A'$ and $||z_2||_{\infty} \leq A''$ and not returning \perp is $(\frac{2A'+1}{2a'+1})^{n+d-1}(\frac{2A''+1}{2a''+1})^d$, which means that the outputs of Trans and Sim have the same distribution.

Trans (sk) Sim (pk) 1: $a \leftrightarrow \mathbb{Z}_q^{< n}[x] := Sam(\rho)$ 2: $y_1 \leftarrow \mathbb{Z}_{\leq a'}^{< n+d-1}[x]$ 1: $a \leftarrow \mathbb{Z}_q^{< n}[x] := Sam(\rho)$ 2: with probability $1 - \left(\frac{2A'+1}{2a'+1}\right)^{n+d-1} \left(\frac{2A''+1}{2a''+1}\right)^d$ 3: $y_2 \leftrightarrow \mathbb{Z}_{\leq a''}^{\leq d}[x]$ 3: output ⊥ 4: $w = a \odot_d y_1 + y_2$ 4: $c \leftarrow D_H$ 5: $c \leftarrow D_H$ 5: $z_1 \leftrightarrow \mathbb{Z}_{\leq A'}^{\leq n+d-1}[x]$ 6: $z_1 = c \odot_{n+d-1} s + y_1$ 6: $z_2 \leftrightarrow \mathbb{Z}_{\leq A''}^{\leq d}[x]$ 7: $z_2 = c \odot_d e + y_2$ 7: output (z_1, z_2, c) 8: if $||z_1||_{\infty} > A'$ or $||z_2||_{\infty} > A''$ then 9: output \perp 10: else 11: output (z_1, z_2, c) 12: end if



Lemma 5.6. The scheme has correctness error $\delta = 1 - (\frac{2A'+1}{2a'+1})^{n+d-1}(\frac{2A''+1}{2a''+1})^d$.

Proof. First, we show that the verification procedure always accepts a honest transcript if $(z_1, z_2) \neq \bot$. Assume that $(z_1, z_2) \neq \bot$. It means that $||z_1||_{\infty} \leq A'$ and $||z_2||_{\infty} \leq A''$. Now we prove that

$$a \odot_d z_1 + z_2 - c \odot_d b = a \odot_d y_1 + y_2.$$

Because of Lemma 4.2, we have that

$$a \odot_d z_1 = a \odot_d (c \odot_{n+d-1} s + y_1)$$
$$= a \odot_d (c \odot_{n+d-1} s) + a \odot_d y_1$$
$$= (a \cdot c) \odot_d s + a \odot_d y_1$$

and

$$c \odot_d b = c \odot_d (a \odot_{d+k} s + e)$$
$$= c \odot_d (a \odot_{d+k} s) + c \odot_d e$$
$$= (c \cdot a) \odot_d s + c \odot_d e.$$

Overall, we obtain:

$$\begin{aligned} a \odot_d z_1 + z_2 - c \odot_d b \\ &= ((a \cdot c) \odot_d s + a \odot_d y_1) + (c \odot_d e + y_2) - ((c \cdot a) \odot_d s + c \odot_d e) \\ &= a \odot_d y_1 + y_2. \end{aligned}$$

Since Sim outputs \perp with the same probability as Trans, we know that the probability to have $(z_1, z_2) = \perp$ is exactly δ .

Lemma 5.7. The identification scheme ID is lossy.

Proof. In the lossy key generation algorithm LossylGen (Figure 5.5), we generate the public key (a, b) uniformly. The public keys generated by IGen and LossylGen are indistinguishable by the MP-LWE assumption. Indeed, for any quantum adversary A against ID, there exists an adversary B trying to distinguish MP-LWE samples from uniform ones such that the loss advantage $\operatorname{Adv}_{\operatorname{ID}}^{loss}(A)$ is equal to the advantage of B.

Lemma 5.8. The identification scheme ID has $d \cdot \log(2a'' + 1)$ bits of min-entropy.

Proof. Indeed, for every commitment ω , we have that:

6

$$\Pr_{a,y_1,y_2}(a \odot_d y_1 + y_2 = \omega) \le \max_{a,y_1} \Pr_{y_2}(y_2 = \omega - a \odot_d y_1) \le \frac{1}{(2a'' + 1)^d}$$

where the first probability is taken over the uniform choice of $a \in \mathbb{Z}_q^{< n}[x]$, $y_1 \in \mathbb{Z}_{\leq a'}^{< n+d-1}[x]$ and $y_2 \in \mathbb{Z}_{\leq a''}^{< d}[x]$. In the second one, the probability is taken over the uniform choice of $y_2 \in \mathbb{Z}_{\leq a''}^{< d}[x]$ and the maximum is taken over all $a \in \mathbb{Z}_q^{< n}[x]$ and $y_1 \in \mathbb{Z}_{\leq a'}^{< n+d-1}[x]$.

LossylGen

1: $\rho \leftrightarrow \{0,1\}^{256}$ 2: $a \leftrightarrow \mathbb{Z}_q^{<n}[x] := Sam(\rho)$ 3: $b \leftrightarrow \mathbb{Z}_q^{<d+k}[x]$ 4: output $\mathsf{pk}_{ls} = (a,b)$

Fig. 5.5: The LossylGen algorithm.

Lemma 5.9. The identification scheme ID is ε_{ls} -lossy-sound, where

$$\varepsilon_{ls} \le \frac{1}{|D_H|} + (4A'+1)^{n+d-1} \cdot (4A''+1)^d \cdot |D_H|^2 \cdot q^{-d}.$$

Proof. We show that relatively to a lossy key pk_{ls} generated by the LossylGen algorithm in Figure 5.5, not even an unbounded quantum adversary can impersonate the prover. This reduces to the computation of the following probability taken over the uniform choice of $a \in \mathbb{Z}_q^{< n}[x]$, $b \in \mathbb{Z}_q^{< d+k}[x]$ and $c \in D_H$:

$$P := \Pr(\exists z_1 \in \mathbb{Z}_{\leq A'}^{< n+d-1}[x], z_2 \in \mathbb{Z}_{\leq A''}^{< d}[x] : a \odot_d z_1 + z_2 - c \odot_d b = w).$$

Let S denote the set of pairs (a, b) such that there exists at most one c for which there exist small z_1, z_2 such that $a \odot_d z_1 + z_2 - c \odot_d b = w$. We can write $P \leq P_1 + P_2$, where

$$P_1 = \Pr((a, b) \in S) \cdot \frac{1}{|D_H|} \le \frac{1}{|D_H|}$$

and

$$P_{2} \leq \Pr((a,b) \notin S) \cdot 1$$

$$\leq \Pr(\exists \ c \neq c', z_{1}, z_{2}, z_{1}', z_{2}' : a \odot_{d} (z_{1} - z_{1}') + z_{2} - z_{2}' - (c - c') \odot_{d} b = 0)$$

$$= \Pr(\exists \ e_{c} \in D_{H} - D_{H} \setminus \{0\}, e_{1} \in \mathbb{Z}_{\leq 2A'}^{\leq n+d-1}, e_{2} \in \mathbb{Z}_{\leq 2A''}^{\leq d} :$$

$$a \odot_{d} e_{1} + e_{2} - e_{c} \odot_{d} b = 0),$$

where a and b are uniformly sampled in $\mathbb{Z}_q^{< n}[x]$, respectively $\mathbb{Z}_q^{< d+k}[x]$, $c, c' \in D_H$, $z_1, z_1 \in \mathbb{Z}_{\leq A'}^{< n+d-1}[x]$, and $z_2, z'_2 \in \mathbb{Z}_{< A''}^{< d}[x]$ and $D_H - D_H$ denotes the set $\{d - d' \mid d, d' \in D_H\}$.

Let us fix $(e_c \neq 0, e_1, e_2)$. The rank of $\mathsf{Toep}(e_c)$ is maximum for $e_c \neq 0$, which means that the function $b \mapsto e_c \odot_d b$ maps an element b from the uniform distribution on $\mathbb{Z}_q^{<d+k}[x]$ to an element b' from the uniform distribution on $\mathbb{Z}_q^{<d}[x]$. We can now write:

$$\Pr(a \odot_d e_1 + e_2 - e_c \odot_d b = 0) = \Pr(b' = a \odot_d e_1 + e_2) = q^{-d},$$

where the first probability is taken over the uniform choice of $a \in \mathbb{Z}_q^{<n}[x]$ and $b \in \mathbb{Z}_q^{<d+k}[x]$ and the second one is taken over the choice of $a \in \mathbb{Z}_q^{<n}[x]$ and $b' \in \mathbb{Z}_q^{<d}[x]$. We conclude that $P_2 \leq (4A'+1)^{n+d-1} \cdot (4A''+1)^d \cdot |D_H|^2 \cdot q^{-d}$.

5.5.2 The signature scheme

In Figure 5.6, we present our digital signature scheme which is obtained by the de-randomized Fiat-Shamir transform of the identification scheme ID. The correctness of the signature scheme follows (see [KLS18, p. 11]) from the correctness of the underlying identification scheme (Lemma 5.6). The scheme is UF-CMA secure in the quantum random oracle model, as discussed in Subsection 2.5.

The signature scheme relies on a hash function $H : \{0,1\}^* \to D_H$, which outputs elements with small norms and will be modelled by a random oracle in the security proof. We refer to [DDLL13] for an efficient method to construct such a hash function.

> Sign (sk = $(s, e, K, \rho), m$) KeyGen 1: $\rho \leftarrow \{0,1\}^{256}$ 1: $a \leftarrow \mathbb{Z}_q^{< n}[x] := Sam(\rho)$ 2: $a \leftarrow \mathbb{Z}_q^{<n}[x] := Sam(\rho)$ 2: i = 03: while $(z_1, z_2) = \perp$ and $i \leq k_l$ do 3: $s \leftrightarrow D_{\mathbb{Z}^{n+d+k-1},\alpha'q}$ 4: $e \leftrightarrow D_{\mathbb{Z}^{d+k},\alpha''q}$ 4: i = i + 15: $y_1 \leftrightarrow \mathbb{Z}_{\leq a'}^{\leq n+d-1}[x] := Sam(K||m||i||0)$ 5: $b = a \odot_{d+k} s + e$ 6: $y_2 \leftrightarrow \mathbb{Z}_{\langle a''}^{\langle d}[x] := Sam(K ||m||i||1)$ 6: $vk = (b, \rho)$ 7: $sk = (s, e, K, \rho)$ 7: $w = a \odot_d y_1 + y_2$ 8: output (sk, vk)8: $c := H(w \parallel m)$ 9: $z_1 = c \odot_{n+d-1} s + y_1$ $z_2 = c \odot_d e + y_2$ 10:if $||z_1||_{\infty} > A'$ or $||z_2||_{\infty} > A''$ then 11: $(z_1, z_2) = \perp$ 12:end if 13:14: end while 15: output (z_1, z_2, c) Verify $(vk = (b, \rho), m, (z_1, z_2, c))$ 1: $a \leftarrow \mathbb{Z}_q^{< n}[x] := Sam(\rho)$ 2: $w = a \odot_d z_1 + z_2 - c \odot_d b$ 3: if $c = H(w||m), ||z_1||_{\infty} \leq A'$ and $||z_2||_{\infty} \leq A''$ then output 1 4:5: else output 0 6: 7: **end if**

Fig. 5.6: The signature scheme.

The key generation algorithm samples $a \leftrightarrow \mathbb{Z}_q^{< n}[x]$ using the extendable function Sam seeded with a 256-bit seed ρ , and then two small secret polynomials $s \leftarrow D_{\mathbb{Z}^{n+d+k-1},\alpha'q}$ and $e \leftarrow D_{\mathbb{Z}^{d+k},\alpha''q}$. It outputs $(b = a \odot_{d+k} s + e, \rho)$ as the verification key vk and (s, e, K, ρ) as the signing key sk, K being a random key for the pseudorandom function $Sam(K \parallel \cdot)$ used in the signature algorithm.

To sign a message m, we first recompute $a \leftrightarrow \mathbb{Z}_q^{<n}[x] := Sam(\rho)$, generate deterministic masking parameters $y_1 \leftrightarrow \mathbb{Z}_{<a'}^{<n+d-1}[x] := Sam(K||m||i||0)$ and $y_2 \leftrightarrow \mathbb{Z}_{<a''}^{<d}[x] := Sam(K||m||i||1)$, where i is the repetition index and compute $w = a \odot_d y_1 + y_2$. Then we compute c := H(w||m), $z_1 = c \odot_{n+d-1} s + y_1$ and $z_2 = c \odot_d e + y_2$. A potential signature is now (z_1, z_2, c) . In order to make the signature pair (z_1, z_2) independent of the signing key, we perform rejection sampling on potential signatures before outputting the right one. A potential signature (z_1, z_2, c) is output if both $||z_1||_{\infty} \leq A'$ and $||z_2||_{\infty} \leq A''$.

To check if (z_1, z_2, c) is a valid signature for a message m, we first recompute $a \leftarrow \mathbb{Z}_q^{< n}[x] := Sam(\rho)$ and $w = a \odot_d z_1 + z_2 - c \odot_d b$ and we accept if $||z_1||_{\infty} \leq A'$, $||z_2||_{\infty} \leq A''$ and c := H(w||m).

5.5.3 Concrete parameters

In this section, we give sample parameters with which our digital signature scheme can be instantiated. The choice of parameters takes into account the correctness error probability, the security and the efficiency of our scheme.

The signing acceptance probability is set to p = 1/3 as in [Lyu16] for a fair comparison. In terms of efficiency, we focus on minimizing the size of a signature. Our signature size is $(n + d - 1) \lceil \log(A') \rceil + d \lceil \log(A'') \rceil + \kappa(\lceil \log(k+1) \rceil + 1) \rceil$ bits. The optimal value of d/n for minimizing the signature length is close to 0.5. As d/n reduces below 0.5, the signature dimension drops. Due to the lossiness condition, d/n and $\log q$ are inversely proportional, so we have to increase n to maintain security, which means that overall the signature length will increase. If d/n increases towards 1, $\log q$ reduces but the signature dimension increases and we cannot reduce the signature length.

The size of our public key (a, b) is $256 + (d + k) \lceil \log(q) \rceil$ bits. Since for our lossiness property in the security proof we need a much larger q than the one used in [Lyu16], our public key becomes larger than the public key used in [Lyu16]. On the other hand, our scheme has significantly shorter signatures. Our savings in MPSign signature length over the scheme in [Lyu16] arise largely from the smaller secret key coordinates in MPSign. As our attack of Section 5.5.5 shows, such savings are not possible in the scheme of [Lyu16] due to the insecurity of PSIS^{\emptyset} with sufficiently small secret coordinates.

In order to set concrete parameters for our scheme achieving λ bits of security, we need to bound from above the advantage of any adversary trying to attack the UF-CMA security of MPSign in the quantum random oracle model by $2^{-\lambda}$. By Theorem 2.3 and Lemma 5.7, it is enough to bound Adv, $\operatorname{Adv}_{\mathsf{PRF}}^{PR}(C)$ and $2^{-d\log(2a'+1)+1}$ by $2^{-\lambda}/5$ and $8(Q_H+1)^2 \cdot \epsilon_{ls}$ by $2^{-\lambda+1}/5$, where the notations are those from Section 5 and Adv stands for the advantage of an adversary trying to solve the MP-LWE_{q,n,d+k,\chi1,\chi2} problem, where both χ_1 and χ_2 are discrete Gaussians of parameters $\alpha''q$, respectively $\alpha'q$. As it is standard in lattice-based cryptography, we further neglect the noise amplification in Theorem 4.3 and assume that the MP-LWE problem with very small secret (with $||s||_{\infty} \approx 1$) is concretely at least as hard as the PLWE^(f) problem with very small secret. Indeed, there are no known attacks on the MP-LWE with small secrets problem that exploit the very small secret when generic algebraic attacks on LWE are protected against (see, e.g., [AG11, ACF⁺15a, ACF⁺15b]). Since the discrete Gaussian distributions of the error and secret have small standard deviation, we assume that we can safely replace them by a corresponding centered binomial distribution, as has been done in many practical lattice-based encryption schemes (see [ADPS16, SSZ19, BDK⁺18], among others).

We use [APS15] in order to estimate both the classical and quantum bit complexities of the primal attack against the $\mathsf{PLWE}^{(f)}$ problem associated to a polynomial f of maximum degree n from the family. The cost models we choose are bkz.sieve for classical security, respectively bkz.qsieve for quantum security.

We present in Table 5.1 a comparison between the efficiency of MPSign and the scheme described

	MPSign	[Lyu16]
public key size	19 KB	9.6 KB
secret key size	$0.7~\mathrm{KB}$	8.8 KB
signature size	13 KB	27 KB
q	$\approx 2^{87}$	$\approx 2^{30}$

	$\lambda_Q = 130$	$\lambda_Q = 89$
n	3800	2500
d	1910	1300
k	512	512
q	$\approx 2^{90.9}$	$\approx 2^{87.3}$
κ	53	53
$ D_H $	$\approx 2^{294}$	$\approx 2^{294}$
$\log A'$	≈ 21.0	≈ 20.4
$\log A''$	≈ 19.4	≈ 18.9
δ	1.004	1.005
$\alpha' q$	$2\sqrt{\pi}$	$2\sqrt{\pi}$
$\alpha''q$	$2\sqrt{\pi}$	$2\sqrt{\pi}$
public key size	26.9 KB	19.5 KB
secret key size	1.06 KB	0.74 KB
signature size	20.1 KB	12.8 KB

Tab. 5.1: Efficiency of MPSign.

Tab. 5.2: Sample parameters for MPSign for λ_Q bits of quantum security.

in [Lyu16]. For the same Hermite factor $\delta_0 = 1.005$ (driving the security level), by choosing n = 2500, d = 1300, k = 512 for our scheme, we manage to shorten the size of a signature by a factor of 2.1. We manage to also shorten the size of the secret key by a factor of 11 at the cost of doubling the size of the public key. Still, one can always only store the seed that is expanded into the secret key during signing.

In the first column of Table 5.2, we provide concrete parameters for MPSign that satisfy both classical and quantum level 1 NIST requirements. Concretely, they achieve $\lambda \ge 143$ for classical adversaries and $\lambda \ge 130$ for quantum adversaries. The second column contains parameters for $\lambda = 89$ bits of quantum security, corresponding to a Hermite factor $\delta = 1.005$.¹

5.5.4 Implementation

We implemented MPSign in Sage (Python) as a proof-of-concept and the source code is publicly available.² For the experiments, we used a MacBook Pro with Intel i7-8559U CPU at 2.7 GHz. Turboboost and hyperthreading were both disabled. For a fair comparison, we also implemented the scheme from [Lyu16]. It is expected that both implementations are slower than if they were implemented with a system language (such as C) with an aim for optimization. Nonetheless, since both implementations use the same Gaussian sampler, the same *hash to challenge* function, and the same polynomial

¹We analyse the $\lambda = 89$ case in order to directly compare with the sample parameters in [Lyu16].

 $^{^{2} \}tt https://github.com/pqc-ntrust/middle-product-LWE-signature$

	[Lyu16]			MPSign		
	min	ave	max	min	ave	max
key generation	22.3	25.9	46.7	14.6	16.3	27.1
signing	111	418	5771	28.3	99.6	713
verification	15.0	30.8	53.0	16.3	18.8	28.6

Tab. 5.3: Performance comparison, in ms.

multiplication algorithm, we believe that the comparison is relatively fair.

We instantiate MPSign and the scheme from [Lyu16] with corresponding parameters achieving $\delta = 1.005$ (for MPSign these parameters may be found in Table 5.2). In both benchmarks we iterated 1000 times, each time with a different seed and a different message to sign. The results of our comparison may be found in Table 5.3. The data are for the average cost in milliseconds. Our scheme is almost twice faster than the one from [Lyu16] in key generation and verification, and four times faster in signing. This is mainly due to the fact that the scheme from [Lyu16] requires scalar multiplications over vectors of polynomials, while our scheme involves a single middle-product (over a somewhat longer polynomial).

5.5.5 An attack on Inhomogeneous $PSIS^{\emptyset}$ with small secrets

In contrast to our hardness result for MP-LWE with small secret coordinates shown in the previous section, here we show a simple efficient attack on the Inhomogeneous $PSIS^{\emptyset}$ problem from [Lyu16] with sufficiently small secret coordinates (such that it has a unique solution). Our algorithm gives a key recovery attack against a small secret variant of the signature scheme of [Lyu16], and shows that a lower bound on the size of the secret key coordinates similar to that in the security proof of [Lyu16] is *necessary* for the security of that signature scheme. MPSign achieves lower signature size than [Lyu16], by using small secret coordinates. The attack presented below shows that a similar improvement in signature size *cannot* be securely achieved in [Lyu16], stressing an MPSign advantage over the approach of [Lyu16].

We now recall the definition of the Inhomogeneous $\mathsf{PSIS}^{\emptyset}$ problem (which we denote by $\mathsf{I}\text{-}\mathsf{PSIS}^{\emptyset}$) from [Lyu16]. The hardness of that problem underlies the security of the key generation algorithm in the signature scheme of [Lyu16]. We note that our definition below is the 'exact' case of the 'approximate' definition in [Lyu16] (with the parameters of [Lyu16, Def. 3.3] set as c = 1, $s = \beta$ and $d_1 = d_2 = d$). This restriction makes our attack even stronger since a solution to the exact problem is also a solution to the 'approximate' problem.

Definition 5.1 (I-PSIS^{\emptyset}). Let n, d > 0. An instance of the I-PSIS $_{q,n,d,k,\beta}^{\emptyset}$ problem consists of a tuple (a_1, \ldots, a_k, t) , where $a_i \leftrightarrow \mathbb{Z}_q^{\leq n}[x]$ for $i = 1, \ldots, k$ and $t = \sum_{i=1}^k a_i \cdot s_i \in \mathbb{Z}_q^{\leq n+d-1}[x]$, where $s_i \leftrightarrow [-\beta, \beta]^{\leq d}[x]$ for $i = 1, \ldots, k$. A solution to the problem is k elements (s'_1, \ldots, s'_k) with $s'_i \in [-\beta, \beta]^{\leq d}[x]$ for $i = 1, \ldots, k$ such that

$$\sum_{i=1}^{k} a_i \cdot s_i' = t$$

Note that the public key of the signature scheme of [Lyu16] consists of an instance of $I-\mathsf{PSIS}^{\emptyset}$, and a solution is a valid secret key.

Our attack on $\text{I-PSIS}^{\emptyset}$ works in the case where s_1, \ldots, s_k is the unique solution, and consists of a simple greedy algorithm that exploits the zero triangles in the Toeplitz matrices associated with the polynomials a_i , to reduce the problem to a sequence of k-dimensional knapsack subproblems: for each

r < d, we recover the k-tuple of coefficients of x^r in the polynomials $s_i(x)$ for i = 1, ..., k. When k is small (as is the case for efficient parameter sets), the attack is efficient.

In more detail, let $t(x) = \sum_{i=1}^{k} a_i(x) \cdot s_i(x) \in \mathbb{Z}_q^{<n+d-1}[x]$ be the target polynomial in an instance of I-PSIS^{\emptyset}. We denote by t_r , $a_{i,r}$ and $s_{i,r}$ the coefficient of x^r in the polynomials $t(x), a_i(x), s_i(x)$, respectively. We observe that for any $r = 0, \ldots, d-1$, the coefficient t_r depends only on the coefficients of x^j for $j \leq r$ of the s_i 's, namely we have

$$t_r = \sum_{i=1}^k \sum_{j=0}^r a_{i,j} \cdot s_{i,r-j} = \sum_{i=1}^k a_{i,0} \cdot s_{i,r} + \sum_{i=1}^k \sum_{j=1}^r a_{i,j} \cdot s_{i,r-j}.$$
 (5.1)

Given an instance (a_1, \ldots, a_k, t) of the I-PSIS $_{q,n,d,k,\beta}^{\emptyset}$ problem, our algorithm works as follows:

- 1 For $r = 0, \ldots, d 1$:
 - (a) Find some vector $s'_{*,r} := (s'_{1,r}, \ldots, s'_{k,r}) \in [-\beta, \beta]^k$ such that

$$t_r = \sum_{i=1}^k a_{i,0} \cdot s'_{i,r} + \sum_{i=1}^k \sum_{j=1}^r a_{i,j} \cdot s'_{i,r-j}.$$
(5.2)

- (b) If no such vector $s'_{*,r}$ exists, return \perp .
- 2 Return (s'_1, \dots, s'_k) , where $s'_i = \sum_{j=0}^{d-1} s'_{i,j} x^j$ for $i = 1, \dots, k$.

Lemma 5.10. Suppose q is prime. With probability $\geq 1 - (4\beta + 1)^k/q$ over the choice of a_1, \ldots, a_k , the solution $(s'_1, \ldots, s'_k) = (s_1, \ldots, s_k)$ to the I-PSIS^{\emptyset}_{q,n,d,k, β} problem is unique, and the above algorithm returns this solution in time $(2\beta + 1)^k \cdot \operatorname{poly}(n, d, \log q)$.

Proof. It follows from (5.1) that the solution $(s'_1, \ldots, s'_k) = (s_1, \ldots, s_k)$ satisfies (5.2) for each r and hence can be output by the algorithm. Now suppose, towards a contradiction, that the algorithm outputs \perp or a different solution $(s'_1, \ldots, s'_k) \neq (s_1, \ldots, s_k)$. Then let $r^* \geq 0$ denote the *least* iteration r of the algorithm where the solution $s'_{*,r^*} := (s'_{1,r^*}, \ldots, s'_{k,r^*})$ to (5.2) for $r = r^*$ is not equal to $s_{*,r^*} := (s_{1,r^*}, \ldots, s_{k,r})$. From (5.2), we have

$$t_{r^*} = \sum_{i=1}^k a_{i,0} \cdot s'_{i,r^*} + \sum_{i=1}^k \sum_{j=1}^r a_{i,j} \cdot s_{i,r^*-j} = \sum_{i=1}^k a_{i,0} \cdot s_{i,r^*} + \sum_{i=1}^k \sum_{j=1}^r a_{i,j} \cdot s_{i,r^*-j},$$

and hence

$$\sum_{i=1}^{k} a_{i,0} \cdot (s_{i,r^*} - s'_{i,r^*}) = 0.$$

As a consequence, the vector $v^* := (s_{1,r^*} - s'_{1,r^*}, \dots, s_{k,r^*} - s'_{k,r^*}) \neq 0$ satisfies $\sum_{i=1}^k a_{i,0}v_i^* = 0$, and $v^* \in [-2\beta, 2\beta]^k$. We claim that such a non-zero vector v^* exists with probability at most $(4\beta + 1)^k/q$ over the uniform choice of the $a_{i,0}$'s. Indeed, since q is prime, the probability that a fixed non-zero vector $v \in [-2\beta, 2\beta]^k$ satisfies $\sum_{i=1}^k a_{i,0}v_i = 0$ is 1/q. A union bound over all $\leq (4\beta + 1)^k$ non-zero vectors in $[-2\beta, 2\beta]^k$ provides the claim. Therefore, the algorithm outputs the unique solution $(s'_1, \dots, s'_k) = (s_1, \dots, s_k)$ with probability at least $1 - (4\beta + 1)^k/q$. The run-time follows since Step 1(a) in the algorithm can be implemented by an exhaustive search through all $(2\beta + 1)^k$ possible values for $s'_{*,r}$.

We observe that the run-time can be reduced to $2^{O(k)} \cdot \operatorname{poly}(n, d, \log q)$ using a lattice closest vector algorithm to solve the k-dimensional knapsack problems.

By Lemma 5.10, our algorithm for $\text{I-PSIS}_{q,n,d,k,\beta}^{\emptyset}$ succeeds with high probability when β is at least slightly smaller than $q^{1/k}/4$, and runs in polynomial time when k = O(1), even for very high degrees nand d. In comparison, the hardness reduction for $\text{I-PSIS}_{q,n,d,k,\beta}^{\emptyset}$ in [Lyu16, Le. 3.4] requires the lower bound $\beta > 2^{\lambda/(kd)-1} \cdot q^{1/k \cdot (1+n/d)}$ (where λ denotes the security parameter and is such that the success probability of the $\text{I-PSIS}^{\emptyset}$ attacker handled by the reduction is $> 2^{-\lambda}$). Our attack gives an efficient key recovery attack against the signature scheme of [Lyu16] with small secrets β . For instance, the recommended parameters of the latter scheme have k = 6 and $q \approx 2^{30}$ and $\beta \approx 2^{11.5}$, but $\beta < 2^3$ will suffice for our attack to succeed. Moreover, heuristically, we expect that our algorithm will succeed with even larger β corresponding to a unique solution. The run-time is likely in practice to be in the order of minutes on a typical laptop ³, using LLL lattice reduction for solving the 6-dimensional knapsack instances; even a brute-force search of each knapsack instance would take in the order of only $(2\beta)^k < 2^{30}$ arithmetic operations. For the above parameters, our LLL-based implementation solved 7 out of 10 (resp. 2 out of 10) instances with $\beta = 7$ (resp. $\beta = 8$), taking about 3 minutes on a 3.1GHz Intel Core i5 CPU.

³https://github.com/pqc-ntrust/middle-product-LWE-signature

CHAPTER 6

OPEN PROBLEMS

We discuss here a list of open problems related to our contributions that we find interesting to be further investigated.

In this thesis, one of the main goals was to explore relationships between PLWE and RLWE. We think that it is interesting to see if we can also relate two different instances of the same problem. Progress into this direction has been done in [BGV12, GHPS12], etc. with applications in fully homomorphic encryption, but their results are restricted to the case of extensions $K \subseteq K'$ of cyclotomic fields. In Chapter 3, we have seen that if we slightly modify the coefficients of a polynomial f to create a new polynomial g, the roots do not change too much. This means that the geometry of $K = \mathbb{Q}[x]/(f)$ and $K' = \mathbb{Q}[x]/(g)$ are very similar. Maybe this similarity in the geometry of K and K' could allow us to link the hardness of the $\mathsf{PLWE}^{(f)}$ problem / the RLWE problem defined using the number field K to the hardness of $\mathsf{PLWE}^{(g)}$ / the RLWE problem defined using the number field K'. Even more, if such connections are possible, one could further try to find a polynomial f for which the $\mathsf{PLWE}^{(f)}$ problem is at least as hard as $\mathsf{PLWE}^{(g)}$ for many polynomials q. This would increase the confidence in choosing that specific f over other polynomials, since any cryptographic system based on $\mathsf{PLWE}^{(f)}$ would remain secure as long as $\mathsf{PLWE}^{(g)}$ remains hard to solve for at least one polynomial g in the respective family. Given our hardness result on MP-LWE from Chapter 4, a completely different strategy to attack the last problem would be to find a polynomial f for which $\mathsf{PLWE}^{(f)}$ is at least as hard as MP-LWE. Finding such an f may be hard. Still, it could be possible to reduce MP-LWE to Module-LWE. Indeed, it can be shown that for an MP-LWE_{q,n+1,n, χ} sample $(a, b = a \odot_n s + e)$, we can write the second component as $b = (a_1s_1 \mod f) + (a_2s_2 \mod g)$, where a_1 and a_2 are two uniformly random polynomials whose coefficients depend on the coefficients of a, s_1 and s_2 are uniformly random and independent polynomials depending on s, $f = x^{n-1} + 1$ and $g = x^{n-1} - 1$. This writing is reminiscent of a Module-LWE sample, but there are two observations to be made: unlike in the Module-LWE case, the polynomials a_1 and a_2 are not independent and there are two different parameterizing modulus f and q involved. For the second difficulty, using a field extension where the roots of both f and g live could be the start of a strategy. An MP-LWE to Module-LWE reduction would be a truly remarkable result. Indeed, combined with our results from Chapter 3 and [AD17], it would imply that MP-LWE is actually equivalent to Module-LWE (and even to ApproxSIVP on module lattices via [LS15]) for many parameterizing polynomials f. In the absence of such an equivalence, MP-LWE seems to be harder than Module-LWE.

Regarding MP-LWE, there are also some other questions that we find stimulating. One of our contributions from Chapter 4 is the hardness proof of MP-LWE with small secrets which works by first reducing PLWE with short secrets to MP-LWE with short secrets and then relying on the reduction from PLWE with uniform secrets to PLWE with short secrets from [ACPS09]. Even if it would not have cryptographic consequences, it would be nice to give an alternative proof, by exhibiting a direct reduction from MP-LWE with uniform secrets to MP-LWE with small secrets, similar to the one which works in the LWE or PLWE cases. Also, in contrast to other algebraic variants of LWE, at the moment

CHAPTER 6. OPEN PROBLEMS

there are few cryptographic primitives whose security relies on MP-LWE. Except for our contributions in Chapter 5, the only primitive based on MP-LWE that we are aware of is the identity-based encryption scheme [LVV19]. It would be interesting to see if [LVV19] can be improved, in the sense of reaching exponential security, but also to design other primitives (e.g. homomorphic encryption) based on the MP-LWE problem.

LIST OF PUBLICATIONS

- [RSSS17] Miruna Rosca, Amin Sakzad, Damien Stehlé and Ron Steinfeld. Middle-Product Learning With Errors. In Proc. of CRYPTO, pages 283-297, 2017. Citations: § 13, 18, 63, 68, and 76.
- [RSW18] Miruna Rosca, Damien Stehlé and Alexandre Wallet. On the Ring-LWE and Polynomial-LWE Problems. In Proc. of EUROCRYPT, pages 146–173, 2018. Citations: § 13 and 18.
- [BDH+20] Shi Bai, Dipayan Das, Ryo Hiromasa, Miruna Rosca, Amin Sakzad, Damien Stehlé, Ron Steinfeld and Zhenfei Zhang. MPSign: A Signature from Small-Secret Middle-Product Learning with Errors. In Proc. of PKC, pages 66-93, 2020. Citations: § 13, 18, 63, and 76.

BIBLIOGRAPHY

- [ABB10] S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In Proc. of EUROCRYPT, pages 553–572. Springer, 2010. Citations: § 10 and 16.
- [ABCP15] M. Abdalla, F. Bourse, A. De Caro, and D. Pointcheval. Simple functional encryption schemes for inner products. In *Proc. of PKC*, page 733–751. Springer, 2015. Citations: § 10 and 16.
- [ACF⁺15a] M. R. Albrecht, C. Cid, J-C. Faugère, R. Fitzpatrick, and L. Perret. Algebraic algorithms for LWE problems. ACM Comm. Computer Algebra, 49(2):62, 2015. Citations: § 89.
- [ACF⁺15b] M. R. Albrecht, C. Cid, J.C. Faugère, R. Fitzpatrick, and L. Perret. On the complexity of the BKW algorithm on LWE. *Designs, Codes and Cryptography*, 74(2):325–354, 2015. Citations: § 89.
- [ACPS09] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Proc. of CRYPTO*, pages 595–618. Springer, 2009. Citations: § 10, 16, 45, and 94.
- [ADPS16] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. Post-quantum key exchange A New Hope. In *Proc. of USENIX*, pages 327–343, 2016. Citations: § 40, 78, 79, and 89.
- [ADRSD15] D. Aggarwal, D. Dadush, O. Regev, and N. Stephens-Davidowitz. Solving the shortest vector problem in 2ⁿ time via discrete gaussian sampling. In *Proc. of STOC*, page 733–742. ACM, 2015. Citations: § 22.
- [AG11] S. Arora and R. Ge. New algorithms for learning in presence of errors. In *Proc. of ICALP*, pages 403–415. Springer, 2011. Citations: § 83 and 89.
- [Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In Proc. of STOC, pages 99–108. ACM, 1996. Citations: § 9, 10, 15, and 16.
- [AKS01] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proc. of STOC*, pages 601–610. ACM, 2001. Citations: § 22.
- [ALS16] S. Agrawal, B. Libert, and D. Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In *Proc. of CRYPTO*, pages 333–362. Springer, 2016. Citations: § 10 and 16.
- [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. Citations: § 89.

- [ASD18] D. Aggarwal and N. Stephens-Davidowitz. Just take the average! an embarrassingly simple 2ⁿ-time algorithm for SVP (and CVP). In Symposium on Simplicity in Algorithms - SOSA, 2018. Citations: § 22.
- [BAA⁺19] N. Bindel, S. Akleylek, E. Alkim, P. S. L. M. Barreto, J. Buchmann, E. Eaton, G. Gutoski, J. Kramer, P. Longa, H. Polat, J. E. Ricardini, and G. Zanon. qTESLA: Algorithm specifications and supporting documentation, NIST PQC round 2 submission document, 2019. http://eprint.iacr.org/2016/461. Citations: § 79.
- [Ban95] W. Banaszczyk. Inequalities for Convex Bodies and Polar Reciprocal Lattices in ℝⁿ. In Discrete and Computational Geometry, volume 13, pages 217–231. Springer, 1995. Citations: § 28.
- [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 *Proc. of ASIACRYPT*, pages 55–81. Springer, 2019. Citations: § 13, 19, 65, 66, and 78.
- [BBdV⁺17] J. Bauch, D. J. Bernstein, H. de Valence, T. Lange, and C. van Vredendaal. Short generators without quantum computers: The case of multiquadratics. In *Proc. of EUROCRYPT*, pages 27–59. Springer, 2017. Citations: § 64.
- [BBPS19] M. Bolboceanu, Z. Brakerski, R. Perlman, and D. Sharma. Order-LWE and the hardness of ring-LWE with entropic secrets. In *Proc. of ASIACRYPT*, pages 91–120. Springer, 2019. Citations: § 13, 19, 44, and 45.
- [BBS16] M. Bellare, B.Poettering, and D. Stebila. From identification to signatures, tightly: A framework and generic transforms. In *Proc. of ASIACRYPT*, pages 435–464. Springer, 2016. Citations: § 35.
- [BCD⁺16] J. Bos, C. Costello, L. Ducas, I. Mironov, M. Naehrig, V. Nikolaenko, A. Raghunathan, and D. Stebila. Frodo: Take off the ring! Practical, quantum-secure key exchange from LWE. In *Proc. of CCS*, pages 1006–1018, 2016. Citations: § 23 and 80.
- [BCLvV16] D. J. Bernstein, C. Chuengsatiansup, T. Lange, and C. van Vredendaal. NTRU Prime, 2016. http://eprint.iacr.org/2016/461. Citations: § 54 and 64.
- [BDE⁺11] J. A. Buchmann, E. Dahmen, S. Ereth, A. Hülsing, and M.Rückert. On the security of the Winternitz one-time signature scheme. In *Proc. of AFRICACRYPT*, pages 363–378. Springer, 2011. Citations: § 35.
- [BDK⁺18] J. W. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, and D. Stehlé. CRYSTALS - Kyber: a CCA-secure module-lattice-based KEM. In *EuroS&P*, pages 353–367, 2018. Citations: § 40, 80, and 89.
- [BFKL93] A. Blum, M. Furst, M. Kearns, and R. Lipton. Cryptographic primitives based on hard learning problems. In Proc. of CRYPTO, pages 278–291. Springer, 1993. Citations: § 23.
- [BG14] S. Bai and S. D. Galbraith. An improved compression technique for signatures based on learning with errors. In *Proc. of CT-RSA*, pages 28–47. Springer, 2014. Citations: § 78.
- [BGV12] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) Fully Homomorphic Encryption without Bootstrapping. In *Innovations in Theoretical Computer Science*, pages 309– 325, 2012. Citations: § 10, 11, 16, 17, and 94.

- [BKW03] A. Blum, A. Kalai, and H. Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM*, 50(4):506–519, 2003. Citations: § 23.
- [BLP⁺13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *Proc. of STOC*, page 575–584. ACM, 2013. Citations: § 10, 16, 24, 42, and 57.
- [BM10] Y. Bugeaud and M. Mignotte. Polynomial root separation. International Journal of Number Theory, 6:587–602, 2010. Citations: § 60.
- [Boy10] X. Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In *Proc. of PKC*, page 499–517. Springer, 2010. Citations: § 77.
- [BS16] J.F. Biasse and F. Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In SODA, page 893–902. Society for Industrial and Applied Mathematics, 2016. Citations: § 64.
- [BV11] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Proc. of FOCS, pages 97–106, 2011. Citations: § 10 and 16.
- [CDPR16] R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. In *Proc. of EUROCRYPT*. Springer, 2016. Citations: § 11, 17, 43, and 64.
- [CDW17] R. Cramer, L. Ducas, and B. Wesolowski. Short Stickelberger class relations and application to Ideal-SVP. In Proc. of EUROCRYPT. Springer, 2017. Citations: § 11, 17, 43, and 64.
- [Che13] K. Cheng. Some complexity results and bit unpredictable for short vector problem, 2013. available at https://eprint.iacr.org/2013/052. Citations: § 22.
- [CHJ⁺02] J-S. Coron, H. Handschuh, M. Joye, P. Paillier, D. Pointcheval, and C. Tymen. A generic chosen-ciphertext secure encryption method. In *Proc. of CT-RSA*, pages 263– 276. Springer, 2002. Citations: § 34.
- [CHKP10] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai Trees, or How to Delegate a Lattice Basis. In Proc. of EUROCRYPT, pages 523–552. Springer, 2010. Citations: § 34 and 77.
- [CIV16a] W. Castryck, I. Iliashenko, and F. Vercauteren. On the tightness of the error bound in Ring-LWE. LMS J. Comput. Math, 2016. Citations: § 42 and 64.
- [CIV16b] W. Castryck, I. Iliashenko, and F. Vercauteren. Provably weak instances of Ring-LWE revisited. In *Proc. of EUROCRYPT*, page 147–167. Springer, 2016. Citations: § 42 and 64.
- [CLS16] H. Chen, K. Lauter, and K. E. Stange. Security considerations for galois non-dual RLWE families. In Proc. of SAC, pages 443–462. Springer, 2016. Citations: § 64.
- [CLS19] H. Chen, K. Lauter, and K. E. Stange. Attacks on search RLWE. SIAM Journal on Applied Algebra and Geometry, 1(1):665–682, 2019. Citations: § 40 and 64.
- [Cona] K. Conrad. The conductor ideal. http://www.math.uconn.edu/~kconrad/blurbs/ gradnumthy/conductor.pdf. Citations: § 44.
- [Conb] K. Conrad. The different ideal. Available at http://www.math.uconn.edu/~kconrad/ blurbs/gradnumthy/different.pdf. Citations: § 47 and 51.

- [Con95] J. B. Conway. Functions of one complex variable. Springer, 1995. Citations: § 52.
- [CZZ18] L. Chen, Z. Zhang, and Z. Zhang. On the hardness of the computational ring-lwr problem and its applications. In *Proc. of ASIACRYPT*, pages 435–464. Springer, 2018. Citations: § 65.
- [DD12] L. Ducas and A. Durmus. Ring-LWE in polynomial rings. In Proc. of PKC, pages 34–51. Springer, 2012. Citations: § 40.
- [DDLL13] L. Ducas, A. Durmus, T. Lepoint, and V. Lyubashevsky. Lattice signatures and bimodal Gaussians. In Proc. of CRYPTO, pages 40–56. Springer, 2013. Citations: § 77 and 88.
- [DFMS19] J. Don, S. Fehr, C. Majenz, and C. Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In *Proc. of CRYPTO*, pages 356–383. Springer, 2019. Citations: § 36 and 79.
- [DH76] W. Diffie and M. Hellman. New directions in cryptography. IEEE transactions on Information Theory, 22(6):644–654, 1976. Citations: § 9 and 15.
- [DKL⁺18] L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS - Dilithium: Sigital signatures from module lattices. In *TCHES*, pages 238–268, 2018. Citations: § 40, 78, and 79.
- [DM14] L. Ducas and D. Micciancio. Improved Short Lattice Signatures in the Standard Model. In Proc. of CRYPTO, pages 335–352. Springer, 2014. Citations: § 34 and 77.
- [dW19] R. de Wolf. Quantum computing: Lecture notes. 2019. available at https://arxiv.org/abs/1907.09415. Citations: § 32.
- [EHL14] K. Eisenträger, S. Hallgren, and K. Lauter. Weak instances of PLWE. In Proc. of SAC, 2014. Citations: § 40 and 64.
- [ELOS15] Y. Elias, K. E. Lauter, E. Ozman, and K. E. Stange. Provably weak instances of Ring-LWE. In Proc. of CRYPTO, pages 63–92. Springer, 2015. Citations: § 64.
- [EU16] E.E.Targhi and D. Unruh. Post-quantum security of the Fujisaki-Okamoto and OAEP transforms. In *Proc. of TCC*, pages 192–216. Springer, 2016. Citations: § 34.
- [FO99] E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In Proc. of CRYPTO, pages 537–554. Springer, 1999. Citations: § 34, 79, and 80.
- [FO13] E. Fujisaki and T. Okamoto. Secure integration of asymmetric and symmetric encryption schemes. *Journal of Cryptology*, 26(1):80–101, 2013. Citations: § 34.
- [FS86] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Proc. of CRYPTO. Springer, 1986. Citations: § 34.
- [GGH96] O. Goldreich, S. Goldwasser, and S. Halevi. Collision-free hashing from lattice problems. Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation, pages 30–39, 1996. Citations: § 10 and 16.
- [GHPS12] C. Gentry, S. Halevi, C. Peikert, and N. P. Smart. Ring switching in BGV-style homomorphic encryption. In Proc. of SCN, pages 19–37, 2012. Citations: § 64 and 94.

- [GLP12] T. Güneysu, V. Lyubashevsky, and T. Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In *Proc. of CHES*, page 499–517. Springer, 2012. Citations: § 77.
- [GMSS99] O. Goldreich, D. Micciancio, S. Safra, and J. Seifert. Approximating shortest lattice vectors is not harder than approximating closest lattice vectors. *Information Processing Letters*, 71(2):55–61, 1999. Citations: § 22.
- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, page 197–206, 2008. Citations: § 10, 16, 23, 28, 76, 77, 78, 80, and 83.
- [GRS08] H. Gilbert, M. Robshaw, and Y. Seurin. HB#: Increasing the Security and Efficiency of HB+. In Proc. of EUROCRYPT, pages 361–378. Springer, 2008. Citations: § 68.
- [GV12] D. Galindo and S. Vivek. A practical leakage-resilient signature scheme in the generic group model. In *Proc. of SAC*, pages 50–65. Springer, 2012. Citations: § 83.
- [HHK17] D. Hofheinz, K. Hövelmannsand, and E. Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Proc. of TCC, pages 341–371. Springer, 2017. Citations: § 34 and 79.
- [HHPW10] J. Hoffstein, N. Howgrave-Graham, J. Pipher, and W. Whyte. Practical lattice-based cryptography: NTRUEncrypt and NTRUSign. In *The LLL Algorithm - Survey and Applications*, pages 349–390. Springer, 2010. Citations: § 41.
- [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. Citations: § 27.
- [Hir18] R. Hiromasa. Digital signatures from the middle-product LWE. In Proc. of ProvSec, pages 239–257. Springer, 2018. Citations: § 78.
- [HPS98] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In Proc. of ANTS, pages 267–288. Springer, 1998. Citations: § 11, 16, and 79.
- [HQZ04] G. Hanrot, M. Quercia, and P. Zimmermann. The middle product algorithm I. Appl. Algebra Engrg. Comm. Comput., 14(6):415–438, 2004. Citations: § 66.
- [HWB17] P. Holzer, T. Wunderer, and J. A. Buchmann. Recovering short generators of principal fractional ideals in cyclotomic fields of conductor $p^{\alpha}q^{\alpha}$. In International Conference in Cryptology in India, page 346–368. Springer, 2017. Citations: § 64.
- [JW05] A. Juels and S.A. Weis. Authenticating Pervasive Devices With Human Protocols. In Proc. of CRYPTO, pages 293–298. Springer, 2005. Citations: § 68.
- [Kan83] R. Kannan. Improved algorithms for integer programming and related lattice problems. In Proc. of STOC, pages 193–206. ACM, 1983. Citations: § 22.
- [Kat10] J. Katz. Digital signatures. Springer, 2010. Citations: § 35 and 78.
- [KL14] J. Katz and Y. Lindell. Introduction to modern cryptography, 2nd edition. CRC Press, 2014. Citations: § 32.
- [KLS18] E. Kiltz, V. Lyubashevsky, and C. Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In *Proc. of EUROCRYPT*, pages 552–586. Springer, 2018. Citations: § 35, 36, 37, 77, 79, 84, and 88.

- [Lam79] L. Lamport. Constructing digital signatures from a one-way function, 1979. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory. Citations: § 35.
- [LM06] V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In Proc. of ICALP, pages 144–155. Springer, 2006. Citations: § 10, 11, 16, 17, 24, 26, 64, 65, and 74.
- [LP11] R. Linden and C. Peikert. Better key sizes (and attacks) for LWE-based encryption. In Proc. of CT-RSA, pages 319–339. Springer, 2011. Citations: § 10 and 16.
- [LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. JACM, 2013, 60(6):43, 2010. Citations: § 11, 16, 17, 40, 41, 42, 45, 46, 47, and 77.
- [LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for Ring-LWE cryptography. In Proc. of EUROCRYPT, pages 35–54. Springer, 2013. Citations: § 56 and 77.
- [LPSS14] S. Ling, D.H. Phan, D. Stehlé, and R. Steinfeld. Hardness of k-LWE and Applications in Traitor Tracing. In Proc. of CRYPTO, pages 315–334. Springer, 2014. Citations: § 28.
- [LS15] A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. Des. Codes Cryptography, 75(3):565–599, 2015. Citations: § 11, 14, 17, 19, 42, 43, 56, 57, 59, 79, and 94.
- [LVV19] A. Lombardi, V. Vaikuntanathan, and T.D. Vuong. Lattice trapdoors and IBE from middle-product LWE. In Proc. of TCC, pages 24–54. Springer, 2019. Citations: § 13, 19, 26, 80, and 95.
- [Lyu08] V. Lyubashevsky. Lattice-based identification schemes secure under active attacks. In Proc. of PKC, pages 162–179. Springer, 2008. Citations: § 10 and 16.
- [Lyu09] V. Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signaturess. In Proc. of ASIACRYPT, page 598–616. Springer, 2009. Citations: § 72 and 77.
- [Lyu12] V. Lyubashevsky. Lattice signatures without trapdoors. In Proc. of EUROCRYPT, page 738–755. Springer, 2012. Citations: § 10, 16, and 77.
- [Lyu16] V. Lyubashevsky. Digital signatures based on the hardness of ideal lattice problems in all rings. In *Proc. of ASIACRYPT*, pages 196–214. Springer, 2016. Citations: § 12, 18, 64, 69, 78, 79, 89, 90, 91, and 93.
- [LZ19] Q. Liu and M. Zhandry. Revisiting post-quantum Fiat-Shamir. In Proc. of CRYPTO, pages 326–355. Springer, 2019. Citations: § 36 and 79.
- [MG02] D. Micciancio and S. Goldwasser. Complexity of Lattice Problems: a cryptographic perspective. Kluwer, 2002. Citations: § 22.
- [Mig00] M. Mignotte. Bounds for the roots of lacunary polynomials. *Journal of Symbolic Computation*, 30(3):325 – 327, 2000. Citations: § 55.
- [MM11] D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *Proc. of CRYPTO*, pages 465–484. Springer, 2011. Citations: § 23.

[MP12] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In Proc. of EUROCRYPT, pages 700-718. Springer, 2012. Citations: § 10, 16, 23, 77, and 80. [MP13] D. Micciancio and C. Peikert. Hardness of SIS and LWE with small parameters. In Proc. of CRYPTO, page 21–39. Springer, 2013. Citations: § 10 and 16. [MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measure. In *Proc. of FOCS*, pages 371–381. IEEE, 2004. Citations: § 10, 16, and 28. [MR09]D. Micciancio and O. Regev. Lattice-based cryptography. In Post-Quantum Cryptography, pages 147–191. Springer, 2009. Citations: § 21. [NIS] NIST. Post-Quantum Cryptography Standardization Process. https://csrc.nist.gov/ Projects/Post-Quantum-Cryptography. Citations: § 15 and 79. [OP01] T. Okamoto and D. Pointcheval. React: Rapid enhanced-security asymmetric cryptosystem transform. In Proc. of CT-RSA, pages 159–175. Springer, 2001. Citations: § 34. [Pan01] V.Y. Pan. Structured matrices and polynomials: Unified superfast algorithms. Springer Science+Business Media New York, 2001. Citations: § 24. [Pei] C. Peikert. Lecture notes of *lattices in cryptography*, taught at the Computer Science and Engineering, University of Michigan. Available at https://web.eecs.umich.edu/ ~cpeikert/. Citations: § 21. [Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In Proc. of STOC, pages 333-342. ACM, 2009. Citations: § 10, 16, 23, and 24. [Pei15] C. Peikert. A decade of lattice-based cryptography, 2015. IACR Cryptology ePrint Archive, report 2015/939. Citations: § 21. [Pei16] C. Peikert. How not to instantiate Ring-LWE. In Proc. of SCN, pages 411–430. Springer, 2016. Citations: § 41, 42, 47, and 64. $[PFH^+19]$ T. Prest, P-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang. Falcon: Algorithm specifications and supporting documentation, 2019. https://falcon-sign.info/falcon.pdf. Citations: § 79. [PHS19] A. Pellet-Mary, G. Hanrot, and D. Stehlé. Approx-SVP in Ideal Lattices with Preprocessing. In Proc. of EUROCRYPT, pages 685–716. Springer, 2019. Citations: § 11, 17, and 64. [PP19] C. Peikert and Z. Pepin. Algebraically structured LWE, revisited. In Proc. of TCC, pages 1-23. Springer, 2019. Citations: § 13, 19, and 66. [PR06] C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Proc. of TCC, pages 145–166. Springer, 2006. Citations: § 10, 11, 16, 17, 28, and 64. [PR07] C. Peikert and A. Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In Proc. of STOC, pages 478–487. Springer, 2007. Citations: § 31. [PRSD17] C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of Ring-LWE for any ring and modulus. In Proc. of STOC, page 461–473. ACM, 2017. Citations: § 11, 12, 14, 17, 19, 41, 42, 43, 56, 59, and 60.

[PVW08]

oblivious transfer. In Proc. of CRYPTO, page 554–571. Springer, 2008. Citations: § 77. [Reg] O. Regev. Lecture notes of lattices in computer science, taught at the Computer Science Tel Aviv University. Available at http://www.cims.nyu.edu/~regev/. Citations: § 21. [Reg05]O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In Proc. of STOC, pages 84–93. ACM, 2005. Citations: § 10, 16, 23, and 24. [Reg09] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. J. ACM, 56(6), 2009. Citations: § 12, 18, 40, 77, 80, and 81. [RSA78] R.L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. Communications of the ACM, 21(2):120-126, 1978. Citations: § 9 and 15. [Sch87] C. P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor.* Comput. Sci., 53(2-3):201-224, 1987. Citations: § 9, 15, and 64. [Sch89] C.P. Schnorr. Efficient identification and signatures for smart cards. In Proc. of CRYPTO, pages 239-252. Springer, 1989. Citations: § 77. [SD15] N. Stephens-Davidowitz. Dimension-preserving reductions between lattice problems. http://noahsd.com/latticeproblems.pdf, 2015. Citations: § 22. Search-to-decision reductions for lattice problems with ap-[SD16] N. Stephens-Davidowitz. proximation factors (slightly) greater than one. APPROX, 2016. available at https: //arxiv.org/pdf/1512.04138.pdf. Citations: § 22. [SE94] C.-P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. Math. Program., 66:181–199, 1994. Citations: § 9 and 15. [Sho94] P. W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In Proceedings of the 35th Annual Symposium on Foundations of Computer Science, pages 124–134. IEEE Computer Society Press, 1994. Citations: § 9 and 15. [Sho99] V. Shoup. Efficient computation of minimal polynomials in algebraic extensions of finite fields. In Proc. of ISSAC, pages 53-58. ACM, 1999. Citations: § 66. [SS11] D. Stehlé and R. Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In Proc. of EUROCRYPT, pages 27–47. Springer, 2011. Citations: § 11, 16, 42, and 56. [SS13] D. Stehlé and R. Steinfeld. Making NTRUEncrypt and NTRUSign as secure standard worst-case problems over ideal lattices, 2013. http://perso.ens-lyon.fr/damien. stehle/NTRU.html. Citations: § 41 and 47. [SSTX09] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In Proc. of ASIACRYPT, pages 617–635. Springer, 2009. Citations: § 10, 11, 16, 17, 40, 46, 64, 65, and 77.

C. Peikert, V. Vaikuntanathan, and B. Waters. A framework for efficient and composable

[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. Citations: § 13, 19, and 79.

- [SSZ19] R. Steinfeld, A. Sakzad, and R. K. Zhao. Practical MP-LWE-based encryption balancing security-risk versus efficiency. *Designs, Codes and Cryptography*, 87(12):2847–2884, 2019. Citations: § 13, 19, 79, and 89.
- [Ste17] P. Stevenhagen. Lecture notes on *number rings*. http://websites.math.leidenuniv. nl/algebra/ant.pdf, 2017. Citations: § 29, 51, 55, and 56.

LIST OF FIGURES

1 Réductions entre les variantes algébriques de LWE. Chaque flèche peut masquer une dégradation du taux d'erreur (et transfert entre rang du module et module dans le cas de [AD17]). Les flèches noires sans références correspondent aux réductions triviales. Les flèches en pointillés correspondent aux résultats présentés dans le Chapitre 3 et les flèches en tirets correspondent aux résultats présentés dans le Chapitre 4. Les flèches verticales en pointillés correspondent à des réductions non-uniformes. Les réductions impliquant PLWE sont analysées pour une famille restreinte de polynômes de définition. 14

1.1	Reductions between algebraic variants of LWE. Each arrow may hide a noise rate degra- dation (and module rank - modulus magnitude transfer in the case of [AD17]). The black arrows without references correspond to trivial reductions. The dotted arrows cor- respond to the results presented in Chapter 3 and the dashed arrows correspond to the results presented in Chapter 4. The vertical dotted arrows correspond to non-uniform re- ductions. The reductions involving PLWE are analyzed for a restricted family of defining																																				
	polynomials.	19																																			
$2.1 \\ 2.2$	A two-dimensional lattice and two of its bases: $\{\mathbf{b}_1, \mathbf{b}_2\}$ and $\{\mathbf{v}_1, \mathbf{v}_2\}$	21																																			
	to $\gamma = 1.5$ approximation factor.	22																																			
2.3	The vector \mathbf{w} is the closest vector to \mathbf{t} belonging to $L = L(\mathbf{v}_1, \mathbf{v}_2)$.	23																																			
2.4	Discrete Gaussian distribution on \mathbb{Z} of standard deviation $r = 1.2$.	29																																			
2.5	The $Pub_{PKE}^{CPA,A}$ experiment.	33																																			
2.6	The $Pub_{PKE}^{CCA,A}$ experiment.	33																																			
2.7	The $UF\text{-}CMA_1$ experiment	35																																			
2.8	The algorithm $Trans(sk)$	36																																			
2.9	The signature SIG obtained via Fiat-Shamir transform from ID. \ldots	37																																			
2.10	The de-randomized signature $DSIG$ obtained via Fiat-Shamir transform from $ID.\ \ldots$.	38																																			
3.1	Relationships between variants of RLWE and PLWE. The dotted box contains the prob- lems studied in this chapter. Each arrow may hide a noise rate degradation (and module rank - modulus magnitude transfer in the case of [AD17]). The top to bottom arrows in the dotted box correspond to non-uniform reductions. The reductions involving PLWE are analyzed for limited family of defining polynomials. The arrows without references correspond to trivial reductions.	43																																			
5.1	PKE scheme from MP-LWE.	81																																			
5.2	Attempt of a dual-Regev PKE scheme from MP-LWE	84																																			
5.3	The identification scheme $(IGen, V, P = (P_1, P_2))$	85																																			
5.4	The transcript $Trans$ and the simulation Sim algorithms	86																																			
5.5	The LossylGen algorithm.	87																																			
5.6	The signature scheme.				•	•		•						•		•		•					•	•			•	•						•	•	88	8
-----	-----------------------	--	--	--	---	---	--	---	--	--	--	--	--	---	--	---	--	---	--	--	--	--	---	---	--	--	---	---	--	--	--	--	--	---	---	----	---
-----	-----------------------	--	--	--	---	---	--	---	--	--	--	--	--	---	--	---	--	---	--	--	--	--	---	---	--	--	---	---	--	--	--	--	--	---	---	----	---

LIST OF TABLES

5.1	Efficiency of MPSign	90
5.2	Sample parameters for MPSign for λ_Q bits of quantum security	90
5.3	Performance comparison, in ms .	90