

PROJECT PROMETHEUS
PRivacy preserving pOst-quantuM systEms
from advanced crypTograpHic mEchanisms
Using lattices



D3.1

Survey on computational problems, cryptanalysis, and basic tools

Contractual submission date
Month 10


Deliverable version
1.0

Actual submission date
May 2019

Main author
Eike Kiltz (RUB)



<http://www.h2020prometheus.eu/>

 h2020prometheus

PROMETHEUS project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 780701. The contents of this deliverable are the responsibility of the PROMETHEUS consortium, and do not necessarily reflect the official views of the European Union.

Document information

Grant agreement no.	780701
Project acronym	PROMETHEUS
Project full title	PRivacy preserving pOst-quantuM systEms from advanced crypTograpHic mEchanisms Using latticeS
Type of action	Research and Innovation Action (RIA)
Topic	H2020-DS-06-2017-Cybersecurity PPP: Cryptography
Project dates	1 st January 2018 (Month 1) / 31 st December 2021 (Month 48)
Duration	48 months
Project URL	http://www.h2020prometheus.eu/
EU Project Officer	Carmen Ifrim
Work package	WP3 – Computational problems, cryptanalysis and basic tools
Deliverable title	Survey on computational problems, cryptanalysis, and basic tools
Deliverable no.	D3.1
Deliverable version	1.0
Deliverable filename	PROMETHEUS-WP3-D3.1.pdf
Nature of deliverable	Report
Dissemination level	Public
Number of pages	29
Responsible partner	RUB (participant number 6)
Authors	Eike Kiltz (RUB), Leo Ducas (CWI) Damien Stehlé (ENSL) Martin Albrecht (RHUL) Tim Güneysu (RUB)

Abstract. This document provides details on the related-work on computational problems, cryptanalysis and lattice trapdoors, as well as a list of open problems.

Keywords: Lattices, Cryptanalysis, Implementations.

Signatures

Written by	Eike Kiltz	RUB	May 2019
Reviewed by	Olivier Sanders	ORA	30/04/2019
Reviewed by	Leo Ducas	CWI	30/04/2019
Approved by	Benoît Libert as Project coordinator	ENSL	07/05/2019
Approved by	Sébastien Canard as Technical leader	ORA	07/05/2019

Partners

ENSL	ENS de Lyon
ORA	Orange SA
CWI	Stiching Centrum Voor Wiskunde En Informatica
IDC	IDC Herzliya
RHUL	Royal Holloway, University of London
RUB	Ruhr-Universität Bochum
SCYTL	Scytl Secure Electronic Voting, S.A.
THA	Thales Communications & Security S.A.S.
TNO	TNO
UPC	Universitat Politècnica de Catalunya · BarcelonaTech
UR1	Université de Rennes 1
WEI	Weizmann Institute of Science

1 Introduction

This document provides a survey on commonly used computational problems on lattices, the state of the art of lattice cryptanalysis, and discusses several important implementation aspects.

Most of lattice-based cryptography relies on the assumption that the SIS and LWE problems (and structured variants thereof like Ring-SIS/LWE or Module-SIS/LWE) are computationally intractable even for quantum computers. Section 2 recalls basic definitions, formally defines the above assumptions, and provides a survey over known techniques to attack them on classical and quantum computers. Section 3 deals with all implementation aspects of lattice-based cryptography, including potential problems with lattice trapdoors and side-channel attacks.

2 Security Issues

We first focus on security issues from a theoretical point of view. After some definitions, we introduce mostly used security assumptions related to lattices, then consider cryptanalysis and finally give some words about quantum random oracles.

2.1 Definition and Notation

Vectors are denoted in bold lower-case letters and bold upper-case letters will denote matrices. The Euclidean and infinity norm of any vector $\mathbf{b} \in \mathbb{R}^m$ will be denoted by $\|\mathbf{b}\|$ and $\|\mathbf{b}\|_\infty$, respectively. The Euclidean norm of matrix $\mathbf{B} \in \mathbb{R}^{m \times n}$ with columns $(\mathbf{b}_i)_{i \leq n}$ is $\|\mathbf{B}\| = \max_{i \leq n} \|\mathbf{b}_i\|$. When \mathbf{B} has full column-rank, we let $\tilde{\mathbf{B}}$ denote its Gram-Schmidt orthogonalization.

2.2 Lattices

A lattice L is the set of integer linear combinations of linearly independent basis vectors $(\mathbf{b}_i)_{i \leq n}$ living in \mathbb{Z}^m . We work with q -ary lattices, for some prime q .

Definition 1 Let $m \geq n \geq 1$, a prime $q \geq 2$ and $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{u} \in \mathbb{Z}_q^n$, define the lattice $\Lambda_q(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m \mid \exists \mathbf{s} \in \mathbb{Z}_q^n \text{ s.t. } \mathbf{A}^\top \cdot \mathbf{s} = \mathbf{e} \bmod q\}$ as well as

$$\Lambda_q^\perp(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m \mid \mathbf{A} \cdot \mathbf{e} = \mathbf{0}^n \bmod q\},$$

as well as

$$\Lambda_q^{\mathbf{u}}(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m \mid \mathbf{A} \cdot \mathbf{e} = \mathbf{u} \bmod q\},$$

which is a shift of the lattice $\Lambda_q^\perp(\mathbf{A})$ since, for any arbitrary $\mathbf{t} \in \Lambda_q^{\mathbf{u}}(\mathbf{A})$, we have $\Lambda_q^{\mathbf{u}}(\mathbf{A}) = \Lambda_q^\perp(\mathbf{A}) + \mathbf{t}$.

For a lattice L , let $\rho_{\sigma, \mathbf{c}}(\mathbf{x}) = \exp(-\pi \|\mathbf{x} - \mathbf{c}\|^2 / \sigma^2)$ for $x \in L$, a vector $\mathbf{c} \in \mathbb{Z}^m$ and a real $\sigma > 0$. The discrete Gaussian of support L , center \mathbf{c} and parameter σ is

$$D_{L, \sigma, \mathbf{c}}(\mathbf{y}) = \rho_{\sigma, \mathbf{c}}(\mathbf{y}) / \rho_{\sigma, \mathbf{c}}(L)$$

for any $\mathbf{y} \in L$, where $\rho_{\sigma, \mathbf{c}}(L) = \sum_{\mathbf{x} \in L} \rho_{\sigma, \mathbf{c}}(\mathbf{x})$. The distribution centered in $\mathbf{c} = \mathbf{0}$ is denoted by $D_{L, \sigma}(\mathbf{y})$.

It is well-known that one can efficiently sample from a Gaussian distribution with lattice support given a sufficiently short basis of the lattice.

Lemma 2.1 ([BLP⁺13, Le. 2.3]) *There exists a PPT algorithm GPVSample that takes as inputs a basis \mathbf{B} of a lattice $L \subseteq \mathbb{Z}^n$ and a rational $\sigma \geq \|\widetilde{\mathbf{B}}\| \cdot \Omega(\sqrt{\log n})$, and outputs vectors $\mathbf{b} \in L$ with distribution $D_{L,\sigma}$.*

We also rely on the trapdoor generation algorithm of Alwen and Peikert [AP09], which refines the technique of Gentry *et al.* [GPV08].

Lemma 2.2 ([AP09, Th. 3.2]) *There is a PPT algorithm TrapGen that takes as inputs $1^n, 1^m$ and an integer $q \geq 2$ with $m \geq \Omega(n \log q)$, and outputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and a basis $\mathbf{T}_\mathbf{A}$ of $\Lambda_q^\perp(\mathbf{A})$ such that \mathbf{A} is within statistical distance $2^{-\Omega(n)}$ to $U(\mathbb{Z}_q^{n \times m})$, and $\|\widetilde{\mathbf{T}_\mathbf{A}}\| \leq \mathcal{O}(\sqrt{n \log q})$.*

2.3 Assumptions and reductions

Most of lattice-based cryptography relies on the assumption that the SIS and LWE problems [Reg05b] (and structured variants thereof like Ring-SIS/LWE [Mic02, SSTX09, LPR10] or Module-SIS/LWE [BGV12, LS15]) are computationally intractable even for quantum computers.

Definition 2 *Let n, m, q, β be functions of $\lambda \in \mathbb{N}$. The Short Integer Solution problem $\text{SIS}_{n,m,q,\beta}$ is, given $\mathbf{A} \leftarrow_s U(\mathbb{Z}_q^{n \times m})$, find $\mathbf{x} \in \Lambda_q^\perp(\mathbf{A})$ with $0 < \|\mathbf{x}\| \leq \beta$.*

The $\text{SIS}_{n,m,q,\beta}^\infty$ is defined in the same way with the difference that the Euclidean norm $\|\mathbf{x}\|$ is replaced by the infinity norm $\|\mathbf{x}\|_\infty$.

If $q \geq \sqrt{n}\beta$ and $m, \beta \leq \text{poly}(n)$, then standard worst-case lattice problems with approximation factors $\gamma = \widetilde{\mathcal{O}}(\beta\sqrt{n})$ reduce to $\text{SIS}_{n,m,q,\beta}$ (see, e.g., [GPV08, Se. 9]).

Definition 3 *Let $n, m \geq 1, q \geq 2$, and let χ be a probability distribution on \mathbb{Z} . For $\mathbf{s} \in \mathbb{Z}_q^n$, let $\mathcal{A}_{\mathbf{s},\chi}$ be the distribution obtained by sampling $\mathbf{a} \leftarrow U(\mathbb{Z}_q^n)$ and $e \leftarrow \chi$, and outputting $(\mathbf{a}, \mathbf{a}^T \cdot \mathbf{s} + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$. The Learning With Errors problem $\text{LWE}_{n,q,\chi}$ asks to distinguish m samples chosen accordingly to $\mathcal{A}_{\mathbf{s},\chi}$ (for $\mathbf{s} \leftarrow U(\mathbb{Z}_q^n)$) and m samples chosen accordingly to $U(\mathbb{Z}_q^n \times \mathbb{Z}_q)$.*

If q is a prime power, $B \geq \sqrt{n}\omega(\log n)$, $\gamma = \widetilde{\mathcal{O}}(nq/B)$, then there exists an efficient sampleable B -bounded distribution χ (i.e., χ outputs samples with norm at most B with overwhelming probability) such that $\text{LWE}_{n,q,\chi}$ is at least as hard as SIVP_γ (see, e.g., [Reg05a, BLP⁺13]). Similarly, if $\alpha q = \Omega(\sqrt{n})$, standard worst-case lattice problems with approximation factors $\gamma = \mathcal{O}(\alpha/n)$ reduce to $\text{LWE}_{n,q,\alpha}$ [Reg05a, BLP⁺13].

Ideal Lattices. Letting q be a prime and $N = 2^r$ for some $r \in \mathbb{N}^+$, we consider the polynomial rings $R = \mathbb{Z}[X]/\langle X^N + 1 \rangle$ and $R_q = \mathbb{Z}_q[X]/\langle X^N + 1 \rangle$. Each ring element $f \in R$ (resp. $f \in R_q$) is thus a polynomial $f = \sum_{i=0}^{N-1} f_i X^i$ of degree at most $N - 1$ in $\mathbb{Z}[X]$ (resp. $\mathbb{Z}_q[X]$). Each $f \in R$ can be associated with the vector $(f_0, f_1, \dots, f_{N-1}) \in \mathbb{Z}^N$ containing its coefficients. When speaking of the norm of a polynomial $f \in R$, we mean the norm of its coefficient vector. We thus use the standard norm definitions $\|f\|_1 = \sum_{i=0}^{N-1} |f_i|$, $\|f\|_2 = (\sum_{i=0}^{N-1} f_i^2)^{1/2}$ and $\|f\|_\infty = \max_i |f_i|$.

For any $g \in R_q$ and $g = \sum_i \bar{g}_i X^i$, we identify each \bar{g}_i with an element $g_i \in [-\frac{q-1}{2}, \frac{q-1}{2}]$ such that $\bar{g}_i = g_i \pmod{q}$. For a positive integer $\alpha > 0$, $S_\alpha = \{a \in R \mid \|a\|_\infty \leq \alpha\}$ denotes the set of all elements in R with ℓ_∞ -norm at most α .

Definition 4 ([LS15]) Let n, m be positive integers and let a real $\beta > 0$. The **Module-SIS** ($M\text{-SIS}_{q,n,m,\beta}$) problem is, given $\mathbf{A} \leftarrow_{\$} U(R_q^{n \times m})$, to find a non-zero $\mathbf{z} \in R$ such that $\mathbf{A} \cdot \mathbf{z} = \mathbf{0}$ and $0 \leq \|\mathbf{z}\| \leq \beta$.

Definition 5 ([LS15]) Let n, m be positive integers and let χ a distribution over R_q . The **Module-LWE** ($M\text{-LWE}_{q,n,m,\chi}$) problem is to distinguish between m uniform samples $(\mathbf{a}_i, b_i) \leftarrow_{\$} U(R_q^n \times R_q)$ and m samples $(\mathbf{a}_i, b_i) \in R_q^n \times R_q$, where $\mathbf{a}_i \leftarrow_{\$} U(R_q^n)$ and $b_i = \mathbf{a}_i^\top \mathbf{s} + e_i$ for each $i \in [m]$, with $\mathbf{s} \leftarrow_{\$} \chi^n$.

Some applications of Module-LWE [DPLS18] use a distribution χ which is simply the uniform distribution over $S_1 = \{a \in R_q \mid \|a\|_\infty \leq 1\}$, in which case M-LWE retains its hardness as long as the number of samples is not too large.

There have been in-depths studies of these assumptions in the last decade, the main focus having been the relationship of (Module-)LWE with standard worst-case lattice problems [Ajt96, Reg05b, LS15]. We will attempt to better understand the quantum hardness of these problems: first with improved quantum reductions between them and second with better relationships with standard quantum problems. As this is pure research, success is highly unpredictable. We cryptanalysis (a hardness proof can be viewed as a “Cryptanalysis with some hints”).

2.4 Cryptanalysis

While the linear algebra formalism offered by LWE and SIS is very convenient for designing and proving the security of cryptographic scheme, their cryptanalysis requires to take a more geometric point of view. Indeed, the SIS problem can be interpreted as the problem of finding a short vector (not necessarily the shortest) in a lattice, while LWE is more of a lattice decoding problem: given a noisy lattice point, separate the noise from the lattice point. This conversion will be explained in Section 2.4.2.

Such problems have a long history, predating lattice-based cryptography. Algorithmically, it starts with the lattice reduction algorithm of LLL [LLL82], but the mathematical notion of lattice reduction itself has been studied for centuries. Nevertheless, the precise behavior of those algorithm remains quite hard to understand precisely: while we have very little doubt about the exponential hardness of lattice problems, determining precisely how hard it is remains a difficult question, and progress is still being made. A brief state of the art is presented in Section 2.4.1.

Furthermore, the growing interest for using lattice with extra algebraic structure to make scheme more efficient (Ring-SIS/LWE) opens the door to other approaches. For quite a while, most structured variants were thought to be essentially as secure as unstructured one. But recent results, using a quantum algorithm for the Hidden-Subgroup problem, have led to demonstrate an asymptotic gap of hardness. Such a gap does not disqualifies structured lattices, but calls for serious scrutiny. Those recent developments are presented in Section 2.4.1.

2.4.1 Lattice Reductions

Lattice reduction algorithms try to produce a basis with short and nearly orthogonal vector, given an integer lattice.

Lattice reduction algorithms have been studied for many years in [LLL82, Sch87, GN08, HPS11, CN11, MW16]. From a theoretical perspective, the best lattice reduction algorithm is the slide reduction algorithm from [GN08]. Alternatively, we may call the

BKZ algorithm [Sch87] and its variants [HPS11, CN11], performing best in practice. Several public implementations of the BKZ algorithm exist [dt17, Sho, AWHT16]. The BKZ algorithm is parameterized by a block size β and consists of repeated calls to an oracle solving the Shortest Vector Problem (SVP) in dimension β combined with calls to the LLL algorithm.

To express the output quality of a lattice reduction, we may relate the shortest vector in the output basis to the volume of the lattice in the *Hermite-factor regime* or to the shortest vector in the lattice, in the *approximation-factor regime*.

The BKZ- β algorithm repeatedly calls an SVP oracle for finding (approximate) shortest vectors in dimension or *block size* β . After BKZ- β reduction, we call the basis *BKZ- β reduced* and in the Hermite-factor regime assume [Che13] that this basis contains a vector of length $\|\mathbf{b}_0\| = \delta_0^d \cdot \text{Vol}(\Lambda)^{1/d}$ where $\delta_0 = \beta^{\Theta(1/\beta)}$ and where $\text{Vol}(\Lambda)$ is the volume of the lattice.

In the approximation-factor regime, we are interested in finding a vector not much longer than some unusually short vector, i.e. a vector that is shorter than predicted for a random lattice. If the unusually short vector is the only vector in that range, this implies we find the target vector. Specifically, we assume that we can find the unusually short vector \mathbf{v} if [ADPS16, AGVW17] if the following condition is satisfied:

$$\sqrt{\beta/d} \|\mathbf{v}\| \leq \delta_0^{2\beta-d} \text{Vol}(\Lambda)^{1/d}. \quad (1)$$

Increasing the parameter β leads to a smaller δ_0 but also leads to an increase in run-time; the run-time grows at least exponentially in β . The two main families of algorithms that can be used to realise the SVP oracle inside BKZ are enumeration and sieving.

Sieving. The cost of sieving on a random lattice of dimension β is $2^{c\beta+o(\beta)}$, where $c = 0.292$ classically [BDGL16]. Some authors replace $o(\beta)$ by a constant based on experiments in [Laa15b], some authors omit it. We note that sieving, while only singly exponential in time, is also exponential in memory. While it was generally assumed that sieving remains uncompetitive in practice compared to enumeration, recent practical improvements suggest that it may soon become competitive [Duc18].

Enumeration. In contrast, enumeration costs $2^{c_1\beta \log \beta + c_2\beta + c_3}$ [Kan83, MW15] but only polynomial memory. The worst case complexity of enumeration is $2^{1/(2e)\beta \log \beta + \Theta(\beta)}$ and curve fitting of data available in [CN11] suggests that this is indeed the cost of BKZ as currently widely implemented [APS15].

2.4.2 Solving LWE

We may pursue one of the two following strategies for solving standard LWE over \mathbb{Z}_q , called the primal and dual strategies. There are other possible strategies but we will point out that they are not competitive.

Primal. Find some \mathbf{s}' such that $\|\mathbf{w} - \mathbf{c}\|$ with $\mathbf{w} = \mathbf{A} \cdot \mathbf{s}'$ is minimized, under the guarantee that \mathbf{w} is not too far from \mathbf{c} . This is known as the Bounded Distance Decoding problem (BDD). To solve BDD, we may embed the BDD instance into a unique SVP (uSVP) instance and apply lattice reduction to solve it. We can then use the estimates for the approximation-factor regime to estimate the required block size. For the primal attack, only one BKZ call suffices, i.e. the attack either succeeds

with high probability or it fails. To solve BDD, we may also perform lattice reduction followed by lattice point enumeration [LP11, LN13].

Dual. Find a short \mathbf{y} in the integral row span of A . This problem is known as the Short Integer Solution problem (SIS). Given such a \mathbf{y} , we can then compute $\langle \mathbf{y}, \mathbf{c} \rangle$. On the one hand, if $\mathbf{c} = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}$, then $\langle \mathbf{y}, \mathbf{c} \rangle = \langle \mathbf{y} \cdot \mathbf{A}, \mathbf{s} \rangle + \langle \mathbf{y}, \mathbf{e} \rangle \equiv \langle \mathbf{y}, \mathbf{e} \rangle \pmod{q}$. If \mathbf{y} is short then $\langle \mathbf{y}, \mathbf{e} \rangle$ is also short. On the other hand, if \mathbf{c} is uniformly random, so is $\langle \mathbf{y}, \mathbf{c} \rangle$, cf., for example, [LP11].

Note that the dual attack solves the decision version of LWE. Thus, applying the Chernoff bound to amplify an advantage ϵ to a constant advantage, we need to perform $\approx 1/\epsilon^2$ experiments and pick by majority vote. However, for the dual attack, too, we can assume that *one* BKZ call is sufficient: A BKZ- β call is followed by $\approx 1/\epsilon^2$ calls to LLL. This assumption is justified heuristically in that we can rerandomise an already reduced basis followed by some light lattice reduction such as LLL to achieve a different basis which is almost as reduced as the input [Alb17]. Alternatively, [ADPS16] argues that sieving outputs exponentially short vectors which can be used (assuming they behave sufficiently close to independent) to amplify the success probability.

Combinatorial Techniques. Both of the above strategies can be augmented with combinatorial techniques. This is often beneficial when the LWE secret follows a small (and/or) sparse distribution, such as a binary $\{0, 1\}$ or ternary $\{-1, 0, 1\}$ distribution. In case of the dual attack, we can view the dual attack as a dimension reduction technique when applied to only a subset of the columns of A [Alb17]. Then, the smaller LWE instance (with a larger error) can be solved e.g. using exhaustive search. In case of the primal attack, we can guess components of the secret to run the attack on a smaller instance.

Other Methods. The dual strategy can also be realized using variants of the BKW algorithm [GJS15, KF15]. However, for the parameter choices considered here, these algorithms are not competitive with lattice-reduction based algorithms.

Furthermore, Arora and Ge proposed an asymptotically efficient algorithm for solving LWE [AG11], which was later improved in [ACF⁺15]. However, these algorithms involve large constants in the exponent, ruling them out for parameters typically considered in cryptography such as here.

2.4.3 Quantum Algorithms

The quantum algorithms can be classified into two classes. The first classes consist of algorithm obtained by accelerating classical algorithm using Grover's quadratic speed-up. The second class consists in polynomial time algorithm based on the Hidden-Subgroup Solvers [EHKS14], a generalization of Shor's quantum factoring algorithm. While there is no indication that general lattices are vulnerable to this second class of algorithms, recent results [BS16, C DPR16, CDW17] have shown a hardness gap between cyclotomic ideal lattices (related to the lattices appearing in Ring-LWE and Ring-SIS) and general lattices.

Grover-like accelerations. Grover's algorithm provides a generic quadratic speed-up to unstructured search problems, and for example forces on to double the key-size

of symmetric cryptographic primitives, assuming a quantum computer can run as fast as a classical one.

It is not straightforward to adapt Grover’s speed-up to lattice algorithm. In the case of lattice enumeration, the work of Montanaro [Mon15] has lead experts to believe in a quadratic speed-up for enumeration, which was very recently confirmed by Aono *et al.* [ANS18]. In the case of Sieving, speed-ups were also found, yet there are much less than quadratic [Laa15a], only decreasing the asymptotic complexity from $2^{.292n+o(n)}$ down to $2^{.265n+o(n)}$. Moreover, such algorithm would require exponential amount of Quantumly accessible RAM, so even if they appear theoretically more promising than enumeration, their practical implementation may be very costly or unfeasible.

This unresolved comparison makes post-quantum security estimate rather delicate, unless one is ready to pay for a comfortable margin. It should also be noted that dust has not yet fully settled, as the latest asymptotic improvements are only a few years old.

Hidden-Subgroup based algorithm for algebraic lattices. The main drawback of lattice-based cryptography is its large memory and bandwidth footprints: a lattice is represented by a basis, i.e. a $m \times n$ matrix for a dimension m of several hundreds. For efficiency reasons, it is tempting to rely on structured lattices, such as lattices generated by a circulant matrix. The earliest example of such a cryptosystem is the NTRUENCRYPT proposal from Hoffstein et al. [HPS98]. Algebraically, those lattices can be viewed as ideals or modules over cyclotomic number fields.

Nevertheless, there is no guarantee that hard lattice problems remain hard on particular classes of structured lattices, and indeed, a series of results [EHKS14, CGS14, BS16, C DPR16, CDW17] have lead to new quantum algorithms solving certain ideal lattice problems. To the best of our knowledge, the same problems remain hard over arbitrary lattices, even with a quantum computer. More precisely, for certain sub-exponential approximation factors α , α -SVP on ideal lattices admit a polynomial-time algorithm, as depicted in Figure 1. A detailed survey on these results can be found in [Duc17].

The impact of these new algorithm remain nevertheless unclear: beyond breaking some exotic lattice-based cryptosystem directly based upon principal ideal of cyclotomic rings (namely, the original FHE scheme of Gentry [Gen09] and the SOLILOQUY encryption scheme [CGS14]), we have not found yet how to apply them to Ring-SIS and Ring-LWE. In other words, while we have proof that such problem are at least as hard as ideal-SVP [Mic02, SSTX09, LPR10], we do not know of any converse reduction.

2.4.4 Open Problems

The most important question is certainly the quantum security of Ring-LWE and Ring-SIS. While not directly affected by the quantum algorithm discussed above, their security must be questioned more than ever, as the best guarantee for a security assumption is the resistance to years of relentless cryptanalysis effort. Such a difficult (and ideally impossible) goal is fortunately paved with interesting intermediate and more realistic results. For example, what can be done using an unlimited pre-computation, or can we determine precisely the performance of the quantum Ideal-SVP algorithm.

Yet another crucial goal when it comes to practice is also to quantify as precisely as possible the hardness under known approaches. To this end, we must go beyond the asymptotic complexity discussed above, and propose heuristic improvements that could




Figure 1: Best known quantum algorithm for general α -SVP (left), and for α -SVP in cyclotomic ideal lattices (right).

be asymptotically negligible, yet quite significant in practice. While such an effort has been done for enumeration, sieving seems to have, up to recent work mostly subject to asymptotic studies; this seems to be a question to prioritize. Developing trade-offs between enumeration and sieving also seems like an important question, as the exponential memory requirement of sieving will become problematic, counterbalancing its better asymptotic running time.

2.5 Quantum Random Oracles

Building provably secure cryptosystems in a classical setting (i.e., without assuming quantum access) has been the focus of intensive research in the last decades. Considering provable security in a post-quantum setting comes with a number of new, sometimes quite subtle, challenges. Tackling these will be one of the goals of this task, which will set the foundations for a number of later tasks.

Provable security means that one proves security of a given cryptosystem relative to some hard computational assumption by means of a security reductions: one shows, via a reduction, that if the computational assumptions holds, then the cryptosystem is secure in some well defined security model. Such a security model is usually given as an interactive probability experiment in which an adversary is run. We refer to standard text books by Goldreich [Gol01] and Katz/Lindell [KL07]. The experiment precisely defines the adversary’s capabilities and its winning condition. In a post-quantum world, an adversary may access some parts of the system via a quantum computer. This is in particular the case for so called “offline” primitives, i.e., cryptographic building blocks that can be implemented and executed without having to interact with the whole system. For example, if the security model provides the adversary access to an idealized hash function, then it is realistic to assume that the access is “quantum”. This is called the quantum random oracle model (QROM) [BDF⁺11]. We will study a number of relevant advanced security models and model their post-quantum counterparts. To this end we will also consider objects like the quantum ideal cipher model or the quantum random permutation model. One central question will be how to even define the concept of “indifferentiability” in the quantum setting.

Once quantum security models are defined, the next step is to provide general techniques for proving security in these models. Again, compared to the classical setting, a number of subtleties arise in the quantum world. For example, one cannot use the rewinding technique, which is a popular technique for proving classical security of signature schemes. Neither can we use pre-image awareness of a quantum random oracle, which has been a crucial technique in almost all classical random oracle proofs.

Open Problems. We intend to develop a toolbox of techniques to deal with general provable security problems in quantum models. This will be crucial input to the other part of the project that will use the toolbox to provide provably secure and efficient cryptosystems in the QROM. We will also look at generic implications between basic primitives in the QROM. Another task will be to find useful relaxations of the QROM such that security proofs are still meaningful in the quantum setting, but are easier to execute. A major problem in the QROM setting is the tightness of the security reduction. In particular, almost all known security reductions are highly non-tight. This, in turn, means a massive increase in the system parameters that usually renders the cryptographic primitive rather inefficient. The central open problem in this setting is to find out if the non-tightness is inherent or if tightness can be improved with more sophisticated security reductions.

3 Implementation issues

We now consider implementation issues related to lattice-based cryptography. We first survey lattice trapdoors and then talk about side channel attacks.

3.1 Trapdoors

A trapdoor function is a function that is efficiently computable by anyone, but can only be efficiently inverted using a secret information, called the trapdoor, built at the time of the function generation. In lattice-based cryptography, the function is typically associated with a lattice, and the trapdoor generally consists of a short basis of that lattice, c.f. Def. 2.2. Unlike trapdoor functions like RSA, those in lattice-based cryptography are not injective. This means that when it is inverted (using the trapdoor), a choice should be made among pre-images. It turned out that the simplest inversion algorithms (e.g., Babai Nearest Plane’s algorithm [Bab86]) lead to a statistical leakage of the trapdoor. Early lattice signature schemes were completely broken because of this leak [NR09, DN12b]. Fortunately, sampling the preimage according to a discrete Gaussian distribution that does not depend on the choice of the trapdoor basis is a provably safe countermeasure [Kle00, GPV08]. While this is in theory a simple algorithm, obtaining a practical implementation raises numerous issues that make the use of trapdoors prohibitive in many cases.

3.1.1 Generating and using lattice trapdoors

Generating lattices with trapdoors has been a major research focus for decades and the first attempt can be traced back to Merkle and Hellman in 1978 [MH06]. Until the NTRU encryption scheme [HPS98], most attempts have been broken, as the generated lattices were more vulnerable than random ones. One of the achievements of modern lattice-based cryptography was to provide a method of generating trapdoors [Ajt99] for lattices that are uniformly distributed in a family for which there is significant evidence that finding short vectors with non-negligible probability is computationally hard [Ajt96]. The most efficient construction to date is the one of Micciancio and Peikert [MP12] that relies on a so-called gadget matrix corresponding to base-2 representation of integers. The construction of [MP12] has other virtues, in particular some form of homomorphism that led to very powerful attribute-based encryption schemes [GVW15], among others.

However, using lattice trapdoors remains quite impractical with current techniques. Two significant obstacles can be pointed out. The first concerns the memory requirement of the technique of [Kle00, GPV08], which requires storing the full Gram-Schmidt-Orthogonalisation (GSO) of the secret short basis. This matrix remains huge even for structured lattices (e.g., module lattices), because the naive GSO algorithm breaks the structure of the basis. Fortunately, new orthogonalisation algorithms should allow to tackle this issue [DP16]. The second obstacle is the concrete arithmetic used for computing the GSO. Exact arithmetic is prohibitive, as the sizes of the numerators and denominators of the rationals involved are too large. Instead, one may rely on floating-point arithmetic, but the known bounds on the numerical precision for ensuring security are currently too high for efficient implementations. Again, recent theoretical results [BLR⁺18] suggest tools to solve this practical issue.

3.1.2 Statistical arguments

The use of statistical arguments is even more prevalent in lattice-based cryptography than in classical cryptography, and this fact is even more acute in the case of algorithms manipulating lattice-trapdoors, for two reasons:

- All known trapdoor samplers are parameterized with a standard deviation, and their correctness (in the sense that they behave similarly to a perfect oracle) increases continuously with the standard deviation. The most natural way to modelise that correct behavior is to use measures of divergence between distributions;
- More specific to trapdoor samplers is the fact that they rely on floating-point arithmetic (FPA): even among lattice-based algorithms, this is quite specific to trapdoor samplers and until now, there has been no way to avoid that in all genericity. FPA needs to be handled with care as it changes the behavior of an algorithm: for example, rounding $0.5 + \epsilon$ to the nearest integer yields very different outcomes whether FPA errors make ϵ to be positive or negative. In the case of lattice trapdoor algorithms, statistical arguments have so far proved to be the best solution for handling uncertainties introduced by FPA;

As usually in cryptography, the statistical distance has been a valuable tool to provide statistical arguments for lattice trapdoor algorithms. This was done in [GPV08] to estimate the required standard deviation, then the analysis was refined and extended to analyze the required precision in [DN12a, LP15]. However, the analysis was not optimal and in particular the required precision (about λ bits of precision, where λ denotes the security parameter) precluded the efficient use of lattice trapdoor algorithms.

More specific divergences can provide tighter arguments as they are more fit to the specificities of lattice-based cryptography, for example the Kullback-Leibler divergence [PDG14, DLP14] allows to decrease the required precision down to about $\lambda/2$ bits.

In lattice-based cryptography, the most powerful metric to this date has been the Rényi divergence [BLL⁺15, Pre17]. As the efficiency of the latter depends on the number of queries granted to an attacker, it allows much tighter proofs under reasonable assumptions for lattice-trapdoor algorithms [BLL⁺15, Pre17, HLS18], but is sometimes delicate to use as it is not a distance.

Other interesting notions have been introduced along the way. The *max-log distance* [MW17] can be combined with the Rényi divergence [Pre17] for maximal

efficiency and ease of use. Finally, [MW18] introduced a new notion of adversarial advantage, which in the case of sampling algorithms allows sharper arguments for decision games. Unfortunately, the arguments are currently not as sharp as in the case of the Rényi divergence for search games.

3.1.3 Implementation

Because of the conjoint use of (high-precision) floating-point arithmetic and discrete Gaussian, it has long been delicate to implement lattice-trapdoor algorithms. Even though the original first mention of trapdoor sampling [GPV08] dates back to 2008, it has not been publicly implemented before the proof-of-concept signature scheme of [EB13].

A line of works by Lyubashevsky, Prest and their co-authors has set to design and implement efficient schemes based on lattice trapdoors. In [DLP14], the IBE from [GPV08] was implemented¹, and algorithmic improvements were proposed and implemented in [LP15]. These two articles provide links to open-source implementations. Another implementation of the IBE from [DLP14] was proposed in [MSO17].

More recently, Ducas and Prest [DP16] proposed and implemented an algorithmic improvement to trapdoor sampling, and their algorithm was subsequently implemented in the signature scheme Falcon² [PFH⁺17] and the hierarchical IBE LATTE [CG17].

In parallel, the lattice trapdoor algorithm of Micciancio and Peikert [MP12], which relies on specific algorithmic tricks, has been implemented in the IBE scheme of [GPR⁺17]. [GM18] proposed a few algorithmic improvements to it, and these improvements were incorporated in the library PALISADE³.

A building block of trapdoor sampling which is particularly tricky to implement is the generation of discrete Gaussians over \mathbb{Z} . Interesting algorithms were proposed and implemented in [MW17, HLS18, KHR⁺18]. Finally, since generation of discrete Gaussians is delicate to implement properly, a testing suite has been proposed⁴ in [HO17] to evaluate the “quality” of Gaussians.

3.1.4 Alternatives to trapdoors

Another avenue for limiting the cost impact of lattice trapdoors is to design solutions that do not necessitate them.

As an alternative to GPV signatures [GPV08] and their extensions (which rely on lattice trapdoors), Lyubashevsky [LM08] proposed a lattice version of Schnorr-like signatures [Sch89]. Nevertheless, a naive adaptation to the lattice setup of Schnorr signatures would also be susceptible to statistical leakages. It is solved by a technique called rejection sampling: knowing two distributions D_1 and D_2 which are somewhat close, it is possible to transform samples following distribution D_1 to samples following D_2 , by rejecting some of them according to a well-crafted rejection process (see, e.g., [Lyu12]). While care is required, this procedure remains nevertheless significantly simpler than generic Gaussian sampling [Kle00, GPV08].

Other cryptographic primitives require only one or a few short vector(s) in the lattice instead of a full short basis. This/these short vector(s) are sometimes called a partial trapdoor. For example, Regev’s public-key encryption scheme uses only

¹<https://github.com/tprest/Lattice-IBE>

²<https://falcon-sign.info/impl/falcon.h.html>

³<https://git.njit.edu/palisade/PALISADE>

⁴<https://github.com/jameshoweee/glitch>

one such vector [Reg05b]. Other examples include the traitor tracing broadcast encryption scheme from [LPSS14] and the inner-product functional encryption scheme from [ALS16].

3.1.5 Open questions

Trapdoor sampling algorithms are not completely practical, for a few reasons.

First, they generally have to rely on floating-point arithmetic, and this brings up the question of the required precision. While recent works [Pre17, PFH⁺17] have shown that the standard value of 53 bits of precision is sufficient in many cases, there remains room for improvement. In particular, these works assume that we are in the presence of search problems, and whether one can achieve the same level of efficiency for decision problems remains an open problem.

An interesting open question which would solve the previous one is whether one can implement lattice trapdoors without floating-point arithmetic. This could be for example achieved via algorithmic tricks.

On a related topic, the generation of discrete Gaussians remains delicate to this date, and no fully satisfying solution has been proposed yet. Therefore, finding such a solution is another open question.

Another interesting direction is to improve the generation of trapdoors themselves. Indeed, the generation of trapdoored lattices leads to quite large lattice parameters impacting the efficiency of the overlying scheme. In particular, it would be interesting to assess whether lattice-coding techniques (e.g., from electrical engineering) could be imported in cryptography for this purpose. Additionally, while the homomorphic properties of lattice trapdoors were mostly a theoretical feature, it could be of practical interest to optimize them for low-degree homomorphisms, and use this limited, low-degree, homomorphisms in cryptographic design.

3.2 Side channel attacks

Implementations of mathematically secure algorithms might still be vulnerable to physical attacks. This applies in particular to embedded applications since the attacker is in possession of the device which executes the algorithm. Such an adversary is free to not only control the input to the device and closely monitor the device, observing its physical properties whilst it performs the cryptographic operations. These physical variables, such as the timings required to perform computations, or the instantaneous power consumed during execution of the algorithm, may be sampled and recorded and used to derive intermediate values of the algorithms. In this report we examine possible attack vectors for implementations of lattice-based algorithms and discuss possible countermeasures to prevent these kind of attacks. We also give an overview of the state of the art in the field of side-channel analysis of lattice-based cryptography.

3.2.1 Physical Attacks

In this section, we discuss attacks exploiting physical properties of an implementation to gain knowledge of the secret key used in the executed algorithm. One distinguishes between passive attacks in which the attacker only monitors information, like execution time, power consumption, or electromagnetic radiation, and active attacks in which the attacker is allowed to interfere in the execution of the cipher.

When dealing with active attacks, one distinguishes between different levels of invasiveness. A non-invasive attacker is only allowed to modify the environment like the temperature, the voltage of the power supply, or the duration of clock cycles. These attacks usually aim to generate a faulty result that can be used to reveal the secret key. A semi-invasive attacker removes the package material of the device and introduces faults by shooting at the a specific location at the device with light or electromagnetic radiation. Invasive attackers aim to even alter the device itself and reverse-engineer the implementation.

Implementations are vulnerable to timing attacks if their execution time depends on secret data. Power analysis exploits the fact that in CMOS technology the dynamic power consumption is dominating in comparison to the static power consumption. An attacker executes the algorithm and measures the power consumption during the execution. The most important types of attacks on the power consumption leakage are simple power analysis (SPA) and differential power analysis (DPA).

Fault attacks The idea of fault attacks is to induce a fault into a circuit and use the faulty output to get information about the secret key. This can be achieved by high temperature, unsupported supply voltage or current, excessively high overclocking, strong electric or magnetic fields, or even ionizing radiation. Fault attacks are usually non-invasive as the induced fault is only temporary and the device is not permanently damaged.

Timing attacks When implementing cryptographic algorithms, the developer has to make sure that the execution time is independent of the secret data that is processed. Otherwise an attacker might be able to exploit the information about the execution time. Such attacks should not only be considered for embedded devices for which the attacker has physical access to, but also remote timing attacks are a threat that must be considered as shown by Brumley and Boneh [BB03]. Timing information can be leaked by conditional branches, instructions with non-constant execution time, and memory accesses that trigger cache hits or misses [Ber05].

Simple power analysis Simple power analysis [KJJR11] works similar to timing attacks. However, while timing attacks exploit the timing information of one or many executions of the algorithms, one or a few power traces of the executed algorithms are used to perform a simple power analysis. An attacker uses visual examination to identify leaking instructions whose execution depends on secret data. Thus, this attack is especially effective when the order of the executed instructions differs from run to run. For instance, an RSA implementation with a naive implementation of the square-and-multiply algorithm can easily be broken by SPA as the square operations and multiply operations are usually easily distinguishable in the power trace. Signal-processing techniques, like frequency filters, might improve the result and make the visual inspection easier.

Differential power analysis While SPA targets the operation-dependency of the power consumption, DPA exploits its data-dependency. Introduced in 1998 by Kocher et al. [KJJR11], DPA (in contrast to SPA) needs many power traces and one analyzes the set of traces with statistical methods. When performing DPA an attacker does not attack the whole key at once, but only a part, e.g. one byte. A DPA is divided in an online phase and an offline phase. During the online phase, the attacker runs a

vast amount of executions of the algorithm to be attacked with different inputs and measures the power consumption of the target device during each run. DPA requires a leakage model that is a prediction of the power consumption. During the offline phase, the attacker guesses the key byte and computes the intermediate value that he considers suitable to apply the power model to. Depending on the power model and the intermediate value, she assigns the corresponding power trace to one of two sets where one contains power traces with high predicted power consumption and one set contains traces with low prediction power consumption. For all power traces, the attacker stores the difference of the means of the sets. If the attack worked, the correct key guess has a much higher difference of means than the other guesses.

t-test A commonly used methodology for side-channel analysis is the t -test leakage detection method initially proposed in [GJJR11, CDG⁺13]. For the non-specific *fixed vs. random* t -test one takes two types of measurements, one with fixed input and one with random input. The t -statistic t is computed as

$$t = \frac{\mu_F - \mu_R}{\sqrt{\frac{\sigma_F^2}{n_F} + \frac{\sigma_R^2}{n_R}}}$$

where μ_F , σ_F^2 , and n_F (resp. μ_R , σ_R^2 , and n_R) denote the mean, variance, and number of measurements set with fixed input (resp. random input). If the value exceeds the threshold $|t| > 4.5$, the test has detected leakage. As this test does not perform an actual attack and does not consider a certain power model it is called *non-specific*. Apart from the *fixed vs. random* t -test it is also possible to perform a *semi-fixed vs. random* t -test. Such a test does not fix the input but some intermediate values, e.g. part of the state of a block cipher to get a more accurate result.

3.2.2 Countermeasures

In this section we discuss different approach to prevent side-channel analysis. Note that there is not a single countermeasure that can be applied to fix all vulnerabilities, in practice usually a combination of countermeasures is applied.

Hiding Hiding countermeasures are applied to raise the difficulty for an attacker to detect sensitive information in a set of power traces. This can be achieved by introducing additional noise or by trying to equalize the power consumption of all operations.

The first approach can be achieved by other computations that are executed in parallel or by shuffling the order of operations. For hardware implementations one can even instantiate dedicated noise generators to randomize the power consumption. If shuffling is applied an attacker needs to perform an extra alignment step before analyzing the power traces. Otherwise the number of required power traces drastically increases.

The second approach is more suitable for hardware implementations as in microcontrollers the developer has only limited influence on the power consumption of an instruction and only one instruction can be executed in parallel (except the microcontroller features SIMD instructions).

Masking The idea behind masking is to split a secret value into several shares. The secret value can only be reconstructed with the knowledge of all shares. The splitting of the secret value can be performed in a Boolean way or in an arithmetic way. Boolean masking means that the XOR-sum of all shares results in the secret value and arithmetic masking means that the arithmetic sum or difference of the shares results in the secret value. There are conversion approaches to switch between arithmetic and Boolean masking [CGTV15]. The major advantage of masking schemes is that they allow to prove the side-channel security of an algorithm. Nevertheless, there are still implementation challenges that have to be taken care of. Otherwise, a provably secure algorithm might still have a side-channel leakage. To achieve higher-order security, it is necessary to split the secret value into more shares.

Constant-time implementation To prevent timing attacks and simple power analysis it is crucial to develop an implementation that has a constant (or at least secret-independent) execution time. Some pitfalls that should be avoided are:

- **Comparison of secret strings:** Such a comparison must not stop at the first unequal character.
- **Branches:** Branches must not be dependent on secret data. Ideally the same branches are taken for every run of the implementation.
- **Table look-ups:** On platforms with a cache, table look-ups can have varying access times. Thus the index must not depend on secret data for such platforms. In some cases it might be necessary to completely disable caches.
- **Compiler optimization:** A developer must take care that the compiler does not remove instructions that are critical for the security of the implementation but irrelevant for its functionality.

Fault countermeasures The most intuitive way to detect a fault is to utilize redundant computations that are used to check the correctness of the result. Spatial redundancy is a possible countermeasure for hardware implementations and means the same operation is executed twice in parallel. This countermeasure has only a small performance overhead but the area consumption doubles. In contrast to that, temporal redundancy means executing another operation after the original operations has been finished. This can either be an additional decryption after an encryption operation to check whether the result matches the original plaintext or simply another encryption to compare both ciphertexts.

For fault attacks that must induce the fault at a specific point in time, it is also possible to randomize the order of the instructions to make an attack harder. Partial reconfiguration on FPGAs can also be used to randomize the location of the circuit that compute the operation. For linear operations error correcting codes can be used to detect faults.

3.2.3 Physical security of lattice-based schemes

In this section, we review the state-of-the-art of research on the physical security of lattice-based cryptography.

Encryption schemes based LWE (and its ring-variant) have also been analyzed for its resistance against side-channel attacks in several works mainly focusing on DPA. The

first approach to secure LWE against DPA has been proposed by Reparaz et al. in 2015 [RRVV15, RRdC⁺16]. The authors attempt to protect the secret key by splitting it into two shares and perform all operations separately on both shares. However, the last step of the algorithm is a decoding function that is not a linear operation and thus requires the knowledge of both shares. To solve this problem [RRVV15] proposed a masked decoder. As this decoder has a number of drawbacks, like being non-deterministic and increasing the failure rate of the scheme, Reparaz et al. [RdCR⁺16] proposed another approach in 2016. In [RdCR⁺16] not the secret key, but the ciphertext is split into two shares. This approach introduces a heavy computational overhead as it requires another run of the decryption during the encryption. Oder et al. [OSPG16] combined the ideas of [RRVV15] and [RdCR⁺16] to avoid the aforementioned problems and also applied a CCA2-conversion to R-LWE to make it secure against adaptive chosen-ciphertext attackers. Furthermore the masking scheme from [OSPG16] has a proof to support its claim. Additionally, [RRVV15, RdCR⁺16, OSPG16] all provide results of practical measurements to demonstrate that the masking schemes indeed prevent a leakage. Oder et al. also discuss the fault sensitivity of R-LWE in [OSPG16]. While these papers focus on encryption and key exchange schemes, the work by Barthe et al. [BBE⁺18] proposes a masking scheme for the GLP signature scheme [GLP12]. Their masking scheme even works for arbitrary orders.

Lattice-based signatures schemes, like BLISS [DDLL13] and GLP [GLP12], have also been analyzed for their vulnerability to fault attacks in [BBK16] and [EFGT16]. Both papers consider instruction-skipping resulting in potential loop aborts and how to exploit such a fault. The work of Bindel et al. [BBK16] furthermore examines the impact of zeroing or randomization of critical values. The proposed countermeasures mainly boil down to redundant computations that are used for correctness checks. Another proposed countermeasure to prevent instruction-skipping is to deliberately induce a segmentation fault by allocating new memory for every intermediate result.

Bruinderink et al. [BHLY16] also found a cache-timing attack on the signature scheme BLISS. More specifically, they attacked the Gaussian sampler that is used to generate noise polynomials in BLISS and are able to extract the secret key with only 3,500 signatures. To prevent timing attacks many implementations of lattice-based schemes provide a constant or secret-independent execution time, like vectorized implementations of the GLP signature scheme and the New Hope key exchange for Intel CPUs [ADPSar, GOPS13]. Furthermore there are also microcontroller implementations of R-LWE that are protected against timing attacks [POG15, OSPG16].

3.2.4 Open problems

The PROMETHEUS project will extend the work on side-channel attacks and countermeasures on implementations of lattice-based cryptography. In particular, we aim to study higher-order masking schemes for lattice-based algorithms as the majority of work from the literature (with the exception of [BBE⁺18]) only covers first-order side-channels. We furthermore want to extend previous work with regard to the submissions to the NIST post-quantum standardization process. Most work on LWE-based schemes can be easily applied to multiple NIST submissions. For signature schemes, it has not been investigated how the analyses conducted on GLP and BLISS can be applied to, for instance, Dilithium. Except for supporting the practical security of NIST submissions, a major target will be targeting side-channel security of advanced constructions, such as identity-based or attribute-based encryption schemes.

4 Conclusion

In the context of lattice-based computational problems, cryptanalysis and basic tools for lattice-based cryptography, a lot of work has already been done and we have considered the most relevant ones in this document. As shown all along this document, it however remains several important open problems and Within WP3 of the PROMETHEUS project, we will consider solving most of them.

References

- [ACF⁺15] Martin R. Albrecht, Carlos Cid, Jean-Charles Faugère, Robert Fitzpatrick, and Ludovic Perret. Algebraic algorithms for LWE problems. *ACM Comm. Computer Algebra*, 49(2):62, 2015. URL: <https://doi.org/10.1145/2815111.2815158>, doi:10.1145/2815111.2815158.
- [ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016.*, pages 327–343. USENIX Association, 2016. URL: <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim>.
- [ADPSar] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange – a new hope. In *Proceedings of the 25th USENIX Security Symposium*. USENIX, (to appear). Document ID: 0462d84a3d34b12b75e8f5e4ca032869, <http://cryptojedi.org/papers/#newhope>.
- [AG11] Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *ICALP 2011, Part I*, volume 6755 of *LNCS*, pages 403–415. Springer, Heidelberg, July 2011.
- [AGVW17] Martin R. Albrecht, Florian Göpfert, Fernando Virdia, and Thomas Wunderer. Revisiting the expected cost of solving uSVP and applications to LWE. In Takagi and Peyrin [TP17], pages 297–322.
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996.
- [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In *Automata, Languages and Programming, 26th International Colloquium, ICALP'99, Prague, Czech Republic, July 11-15, 1999, Proceedings*, volume 1644 of *Lecture Notes in Computer Science*, pages 1–9. Springer, 1999. URL: https://doi.org/10.1007/3-540-48523-6_1, doi:10.1007/3-540-48523-6_1.
- [Alb17] Martin R. Albrecht. On dual lattice attacks against small-secret LWE and parameter choices in HELib and SEAL. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 103–129. Springer, Heidelberg, April / May 2017.
- [ALS16] Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III*, volume 9816 of *Lecture Notes in Computer Science*, pages 333–362. Springer, 2016. URL: https://doi.org/10.1007/978-3-662-53015-3_12, doi:10.1007/978-3-662-53015-3_12.

- [ANS18] Yoshinori Aono, Phong Q. Nguyen, and Yixin Shen. Quantum lattice enumeration and tweaking discrete pruning. *Cryptology ePrint Archive, Report 2018/546*, 2018. <https://eprint.iacr.org/2018/546>.
- [AP09] J. Alwen and C. Peikert. Generating shorter bases for hard random lattices. In *STACS 2009*, volume 3 of *LIPICs*, pages 75–86. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2009.
- [APS15] Martin R Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.
- [AWHT16] Yoshinori Aono, Yuntao Wang, Takuya Hayashi, and Tsuyoshi Takagi. Improved progressive BKZ algorithms and their precise cost estimation by sharp simulator. In Fischlin and Coron [FC16], pages 789–819. doi: 10.1007/978-3-662-49890-3_30.
- [Bab86] László Babai. On lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. URL: <https://doi.org/10.1007/BF02579403>, doi:10.1007/BF02579403.
- [BB03] David Brumley and Dan Boneh. Remote timing attacks are practical. In *Proceedings of the 12th USENIX Security Symposium, Washington, D.C., USA, August 4-8, 2003*, 2003. URL: <https://www.usenix.org/conference/12th-usenix-security-symposium/remote-timing-attacks-are-practical>.
- [BBE⁺18] Gilles Barthe, Sonia Belaïd, Thomas Espitau, Pierre-Alain Fouque, Benjamin Grégoire, Mélissa Rossi, and Mehdi Tibouchi. Masking the glp lattice-based signature scheme at any order. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018*, pages 354–384, Cham, 2018. Springer International Publishing.
- [BBK16] Nina Bindel, Johannes Buchmann, and Juliane Krämer. Lattice-based signature schemes and their sensitivity to fault attacks. In *Fault Diagnosis and Tolerance in Cryptography (FDTC), 2016 Workshop on*, pages 63–77. IEEE, 2016.
- [BDF⁺11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Lee and Wang [LW11], pages 41–69.
- [BDGL16] Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. In Krauthgamer [Kra16], pages 10–24. doi:10.1137/1.9781611974331.ch2.
- [Ber05] Daniel J Bernstein. Cache-timing attacks on aes, 2005.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012*, pages 309–325. ACM, January 2012.

- [BHLY16] Leon Groot Bruinderink, Andreas Hülsing, Tanja Lange, and Yuval Yarom. Flush, gauss, and reload—a cache attack on the bliss lattice-based signature scheme. In *International Conference on Cryptographic Hardware and Embedded Systems*, pages 323–345. Springer, 2016.
- [BLL⁺15] Shi Bai, Adeline Langlois, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 3–24. Springer, Heidelberg, November / December 2015. doi : 10 . 1007/978-3-662-48797-6_1.
- [BLP⁺13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. On the classical hardness of learning with errors. In *STOC*, 2013.
- [BLR⁺18] Shi Bai, Tancrède Lepoint, Adeline Roux-Langlois, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the rényi divergence rather than the statistical distance. *J. Cryptology*, 31(2):610–640, 2018. URL: <https://doi.org/10.1007/s00145-017-9265-9>, doi : 10 . 1007/s00145-017-9265-9.
- [BS16] Jean-François Biasse and Fang Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In Krauthgamer [Kra16], pages 893–902. doi : 10 . 1137/1 . 9781611974331 . ch64.
- [CDG⁺13] Jeremy Cooper, Elke Demulder, Gilbert Goodwill, Joshua Jaffe, Gary Kenworthy, and Pankaj Rohatgi. Test Vector Leakage Assessment (TVLA) Methodology in Practice. International Cryptographic Module Conference, 2013.
- [CDPR16] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 559–585. Springer, Heidelberg, May 2016. doi : 10 . 1007/978-3-662-49896-5_20.
- [CDW17] Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-SVP. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 324–348. Springer, Heidelberg, April / May 2017.
- [CG17] Peter Campbell and Michael Groves. Practical post-quantum hierarchical identity-based encryption. 16th IMA International Conference on Cryptography and Coding, 2017. <http://www.qub.ac.uk/sites/CSIT/FileStore/Filetoupload,785752,en.pdf>.
- [CGS14] Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: A cautionary tale. In *ETSI 2nd Quantum-Safe Crypto Workshop*, pages 1–9, 2014.
- [CGTV15] Jean-Sébastien Coron, Johann Großschädl, Mehdi Tibouchi, and Praveen Kumar Vadnala. Conversion from arithmetic to boolean masking

- with logarithmic complexity. In *International Workshop on Fast Software Encryption*, pages 130–149. Springer, 2015.
- [Che13] Yuanmi Chen. *Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe*. PhD thesis, Paris 7, 2013.
- [CN11] Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better lattice security estimates. In Lee and Wang [LW11], pages 1–20.
- [DDLL13] Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In *Advances in Cryptology—CRYPTO 2013*, pages 40–56. Springer, 2013.
- [DLP14] Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 22–41. Springer, Heidelberg, December 2014. doi : 10.1007/978-3-662-45608-8_2.
- [DN12a] Léo Ducas and Phong Q. Nguyen. Faster Gaussian lattice sampling using lazy floating-point arithmetic. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 415–432. Springer, Heidelberg, December 2012. doi : 10.1007/978-3-642-34961-4_26.
- [DN12b] Léo Ducas and Phong Q. Nguyen. Learning a zonotope and more: Cryptanalysis of ntrusign countermeasures. In *Advances in Cryptology - ASIACRYPT 2012 - 18th International Conference on the Theory and Application of Cryptology and Information Security, Beijing, China, December 2-6, 2012. Proceedings*, volume 7658 of *Lecture Notes in Computer Science*, pages 433–450. Springer, 2012. URL: https://doi.org/10.1007/978-3-642-34961-4_27, doi : 10.1007/978-3-642-34961-4_27.
- [DP16] Léo Ducas and Thomas Prest. Fast fourier orthogonalization. In *Proceedings of the ACM on International Symposium on Symbolic and Algebraic Computation, ISSAC 2016, Waterloo, ON, Canada, July 19-22, 2016*, pages 191–198, 2016. URL: <http://doi.acm.org/10.1145/2930889.2930923>, doi : 10.1145/2930889.2930923.
- [DPLS18] R. Del Pino, V. Lyubashevsky, and G. Seiler. Lattice-based group signatures and zero-knowledge proofs of automorphism stability. In *ACM-CCS*, 2018.
- [dt17] The FPLLL development team. FPLLL, a lattice reduction library. Available at <https://github.com/fplll/fplll>, 2017. URL: <https://github.com/fplll/fplll>.
- [Duc17] Léo Ducas. Advances on quantum cryptanalysis of ideal lattices. *Nieuw Archief voor Wiskunde*, 5:184–189, 2017.
- [Duc18] Léo Ducas. Shortest vector from lattice sieving: A few dimensions for free. In Nielsen and Rijmen [NR18], pages 125–145. doi : 10.1007/978-3-319-78381-9_5.

- [EB13] Rachid El Bansarkhani and Johannes Buchmann. Improvement and efficient implementation of a lattice-based signature scheme. *Cryptology ePrint Archive*, Report 2013/297, 2013. <http://eprint.iacr.org/2013/297>.
- [EFGT16] Thomas Espitau, Pierre-Alain Fouque, Benoît Gérard, and Mehdi Tibouchi. Loop-abort faults on lattice-based fiat–shamir and hash-and-sign signatures. *Cryptology ePrint Archive*, Report 2016/449, 2016. <http://eprint.iacr.org/2016/449>.
- [EHKS14] Kirsten Eisenträger, Sean Hallgren, Alexei Kitaev, and Fang Song. A quantum algorithm for computing the unit group of an arbitrary degree number field. In David B. Shmoys, editor, *46th ACM STOC*, pages 293–302. ACM Press, May / June 2014.
- [FC16] Marc Fischlin and Jean-Sébastien Coron, editors. *EUROCRYPT 2016, Part I*, volume 9665 of LNCS. Springer, Heidelberg, May 2016.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
- [GJJR11] G. Goodwill, B. Jun, J. Jaffe, and P. Rohatgi. A testing methodology for side channel resistance validation. In *NIST non-invasive attack testing workshop*, 2011.
- [GJS15] Qian Guo, Thomas Johansson, and Paul Stankovski. Coded-BKW: Solving LWE using lattice codes. In Gennaro and Robshaw [GR15], pages 23–42. doi: 10.1007/978-3-662-47989-6_2.
- [GLP12] Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, pages 530–547, 2012. URL: http://dx.doi.org/10.1007/978-3-642-33027-8_31, doi: 10.1007/978-3-642-33027-8_31.
- [GM18] Nicholas Genise and Daniele Micciancio. Faster gaussian sampling for trapdoor lattices with arbitrary modulus. In Nielsen and Rijmen [NR18], pages 174–203. doi: 10.1007/978-3-319-78381-9_7.
- [GN08] Nicolas Gama and Phong Q. Nguyen. Finding short lattice vectors within Mordell’s inequality. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 207–216. ACM Press, May 2008.
- [Gol01] Oded Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, Cambridge, UK, 2001.
- [GOPS13] Tim Güneysu, Tobias Oder, Thomas Pöppelmann, and Peter Schwabe. Software speed records for lattice-based signatures. In *Post-Quantum Cryptography - 5th International Workshop, PQCrypto 2013, Limoges, France, June 4-7, 2013. Proceedings*, pages 67–82, 2013. URL: http://dx.doi.org/10.1007/978-3-642-38616-9_5, doi: 10.1007/978-3-642-38616-9_5.

- [GPR⁺17] Kamil Doruk Gür, Yuriy Polyakov, Kurt Rohloff, Gerard W. Ryan, and Erkey Savaş. Implementation and evaluation of improved gaussian sampling for lattice trapdoors. *Cryptology ePrint Archive*, Report 2017/285, 2017. <http://eprint.iacr.org/2017/285>.
- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
- [GR15] Rosario Gennaro and Matthew J. B. Robshaw, editors. *CRYPTO 2015, Part I*, volume 9215 of *LNCS*. Springer, Heidelberg, August 2015.
- [GVW15] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. *J. ACM*, 62(6):45:1–45:33, 2015. URL: <http://doi.acm.org/10.1145/2824233>, doi:10.1145/2824233.
- [HLS18] Andreas Hülsing, Tanja Lange, and Kit Smeets. Rounded gaussians - fast and secure constant-time sampling for lattice-based crypto. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part II*, volume 10770 of *LNCS*, pages 728–757. Springer, Heidelberg, March 2018.
- [HO17] James Howe and Máire O’Neill. GLITCH: A discrete gaussian testing suite for lattice-based cryptography. *Cryptology ePrint Archive*, Report 2017/438, 2017. <http://eprint.iacr.org/2017/438>.
- [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *Algorithmic Number Theory, Third International Symposium, ANTS-III, Portland, Oregon, USA, June 21-25, 1998, Proceedings*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998. URL: <https://doi.org/10.1007/BFb0054868>, doi:10.1007/BFb0054868.
- [HPS11] Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 447–464. Springer, Heidelberg, August 2011.
- [Kan83] Ravi Kannan. Improved algorithms for integer programming and related lattice problems. In *15th ACM STOC*, pages 193–206. ACM Press, April 1983.
- [KF15] Paul Kirchner and Pierre-Alain Fouque. An improved BKW algorithm for LWE with applications to cryptography and lattices. In Gennaro and Robshaw [GR15], pages 43–62. doi:10.1007/978-3-662-47989-6_3.
- [KHR⁺18] Ayesha Khalid, James Howe, Ciara Rafferty, Francesco Regazzoni, and Máire O’Neill. Compact, scalable, and efficient discrete gaussian samplers for lattice-based cryptography. *IACR Cryptology ePrint Archive*, 2018:265, 2018.
- [KJJR11] Paul Kocher, Joshua Jaffe, Benjamin Jun, and Pankaj Rohatgi. Introduction to differential power analysis. *Journal of Cryptographic Engineering*, 1(1):5–27, 2011.

- [KL07] Jonathan Katz and Juhuda Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC Press, aug 2007.
- [Kle00] Philip N. Klein. Finding the closest lattice vector when it's unusually close. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, January 9-11, 2000, San Francisco, CA, USA.*, pages 937–941. ACM/SIAM, 2000. URL: <http://dl.acm.org/citation.cfm?id=338219.338661>.
- [Kra16] Robert Krauthgamer, editor. *27th SODA*. ACM-SIAM, January 2016.
- [Laa15a] Thijs Laarhoven. *Search problems in cryptography: From fingerprinting to lattice sieving*. PhD thesis, Eindhoven University of Technology, 2015.
- [Laa15b] Thijs Laarhoven. Sieving for shortest vectors in lattices using angular locality-sensitive hashing. In Gennaro and Robshaw [GR15], pages 3–22. doi : 10.1007/978-3-662-47989-6_1.
- [LLL82] A.K. Lenstra, Jr. Lenstra, H.W., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982. doi : 10.1007/BF01457454.
- [LM08] Vadim Lyubashevsky and Daniele Micciancio. Asymptotically efficient lattice-based digital signatures. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2008, New York, USA, March 19-21, 2008.*, volume 4948 of *Lecture Notes in Computer Science*, pages 37–54. Springer, 2008. URL: https://doi.org/10.1007/978-3-540-78524-8_3, doi : 10.1007/978-3-540-78524-8_3.
- [LN13] Mingjie Liu and Phong Q. Nguyen. Solving BDD by enumeration: An update. In Ed Dawson, editor, *CT-RSA 2013*, volume 7779 of *LNCS*, pages 293–309. Springer, Heidelberg, February / March 2013. doi : 10.1007/978-3-642-36095-4_19.
- [LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for LWE-based encryption. In Aggelos Kiayias, editor, *CT-RSA 2011*, volume 6558 of *LNCS*, pages 319–339. Springer, Heidelberg, February 2011.
- [LP15] Vadim Lyubashevsky and Thomas Prest. Quadratic time, linear space algorithms for Gram-Schmidt orthogonalization and Gaussian sampling in structured lattices. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 789–815. Springer, Heidelberg, April 2015. doi : 10.1007/978-3-662-46800-5_30.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Heidelberg, May / June 2010.
- [LPSS14] San Ling, Duong Hieu Phan, Damien Stehlé, and Ron Steinfeld. Hardness of k-LWE and applications in traitor tracing. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 315–334. Springer, Heidelberg, August 2014. doi : 10.1007/978-3-662-44371-2_18.

- [LS15] A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3), 2015.
- [LW11] Dong Hoon Lee and Xiaoyun Wang, editors. *ASIACRYPT 2011*, volume 7073 of *LNCS*. Springer, Heidelberg, December 2011.
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755. Springer, 2012. URL: https://doi.org/10.1007/978-3-642-29011-4_43, doi: 10.1007/978-3-642-29011-4_43.
- [MH06] R. Merkle and M. Hellman. Hiding information and signatures in trapdoor knapsacks. *IEEE Trans. Inf. Theor.*, 24(5):525–530, September 2006. URL: <https://doi.org/10.1109/TIT.1978.1055927>, doi: 10.1109/TIT.1978.1055927.
- [Mic02] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *43rd FOCS*, pages 356–365. IEEE Computer Society Press, November 2002.
- [Mon15] Ashley Montanaro. Quantum walk speedup of backtracking algorithms. *arXiv preprint arXiv:1509.02374*, 2015.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012.
- [MSO17] Sarah McCarthy, Neil Smyth, and Elizabeth O’Sullivan. A practical implementation of identity-based encryption over NTRU lattices. In Máire O’Neill, editor, *16th IMA International Conference on Cryptography and Coding*, volume 10655 of *LNCS*, pages 227–246. Springer, Heidelberg, December 2017.
- [MW15] Daniele Micciancio and Michael Walter. Fast lattice point enumeration with minimal overhead. In Piotr Indyk, editor, *26th SODA*, pages 276–294. ACM-SIAM, January 2015. doi: 10.1137/1.9781611973730.21.
- [MW16] Daniele Micciancio and Michael Walter. Practical, predictable lattice basis reduction. In Fischlin and Coron [FC16], pages 820–849. doi: 10.1007/978-3-662-49890-3_31.
- [MW17] Daniele Micciancio and Michael Walter. Gaussian sampling over the integers: Efficient, generic, constant-time. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 455–485. Springer, Heidelberg, August 2017.
- [MW18] Daniele Micciancio and Michael Walter. On the bit security of cryptographic primitives. In Nielsen and Rijmen [NR18], pages 3–28. doi: 10.1007/978-3-319-78381-9_1.

- [NR09] Phong Q. Nguyen and Oded Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *Journal of Cryptology*, 22(2):139–160, April 2009.
- [NR18] Jesper Buus Nielsen and Vincent Rijmen, editors. *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*. Springer, Heidelberg, April / May 2018.
- [OSPG16] Tobias Oder, Tobias Schneider, Thomas Pöppelmann, and Tim Güneysu. Practical cca2-secure and masked ring-lwe implementation. *IACR Cryptology ePrint Archive*, 2016:1109, 2016. URL: <http://eprint.iacr.org/2016/1109>.
- [PDG14] Thomas Pöppelmann, Léo Ducas, and Tim Güneysu. Enhanced lattice-based signatures on reconfigurable hardware. In Lejla Batina and Matthew Robshaw, editors, *CHES 2014*, volume 8731 of *LNCS*, pages 353–370. Springer, Heidelberg, September 2014. doi:10.1007/978-3-662-44709-3_20.
- [PFH⁺17] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon. Technical report, National Institute of Standards and Technology, 2017. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions>.
- [POG15] Thomas Pöppelmann, Tobias Oder, and Tim Güneysu. High-performance ideal lattice-based cryptography on 8-bit ATxmega microcontrollers. In *Progress in Cryptology - LATINCRYPT 2015 - 4th International Conference on Cryptology and Information Security in Latin America, Guadalajara, Mexico, August 23-26, 2015, Proceedings*, pages 346–365, 2015. URL: http://dx.doi.org/10.1007/978-3-319-22174-8_19, doi:10.1007/978-3-319-22174-8_19.
- [Pre17] Thomas Prest. Sharper bounds in lattice-based cryptography using the Rényi divergence. In Takagi and Peyrin [TP17], pages 347–374.
- [RdCR⁺16] Oscar Reparaz, Ruan de Clercq, Sujoy Sinha Roy, Frederik Vercauteren, and Ingrid Verbauwhede. Additively homomorphic ring-lwe masking. In *Post-Quantum Cryptography - 7th International Workshop, PQCrypto 2016, Fukuoka, Japan, February 24-26, 2016, Proceedings*, pages 233–244, 2016. URL: http://dx.doi.org/10.1007/978-3-319-29360-8_15, doi:10.1007/978-3-319-29360-8_15.
- [Reg05a] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.
- [Reg05b] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.
- [RRdC⁺16] Oscar Reparaz, Sujoy Sinha Roy, Ruan de Clercq, Frederik Vercauteren, and Ingrid Verbauwhede. Masking ring-lwe. *J. Cryptographic Engineering*, 6(2):139–153, 2016. URL: <http://dx.doi.org/10.1007/s13389-016-0126-5>, doi:10.1007/s13389-016-0126-5.

- [RRVV15] Oscar Reparaz, Sujoy Sinha Roy, Frederik Vercauteren, and Ingrid Verbauwhede. A masked ring-lwe implementation. In *Cryptographic Hardware and Embedded Systems - CHES 2015 - 17th International Workshop, Saint-Malo, France, September 13-16, 2015, Proceedings*, pages 683–702, 2015. URL: http://dx.doi.org/10.1007/978-3-662-48324-4_34, doi:10.1007/978-3-662-48324-4_34.
- [Sch87] Claus-Peter Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.
- [Sch89] Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 239–252. Springer, 1989. URL: https://doi.org/10.1007/0-387-34805-0_22, doi:10.1007/0-387-34805-0_22.
- [Sho] Victor Shoup. Number theory library 5.5.2 (ntl) for c++. <http://www.shoup.net/ntl/>.
- [SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 617–635. Springer, Heidelberg, December 2009.
- [TP17] Tsuyoshi Takagi and Thomas Peyrin, editors. *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*. Springer, Heidelberg, December 2017.