

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



D4.2

Intermediate results on building blocks for practical advanced protocols

Contractual submission date
Month 24


Deliverable version
1.0

Actual submission date
December 2019

Main author
Devika Sharma and Zvika Brakerski
(WEI)



<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	WP4 – Building blocks for practical advanced protocols
Deliverable title	Intermediate results on building blocks for practical advanced protocols
Deliverable no.	D4.2
Deliverable version	1.0
Deliverable filename	PROMETHEUS-780701-WP4-D4.2.pdf
Nature of deliverable	Report
Dissemination level	Public
Number of pages	47
Responsible partner	WEI (participant number 12)
Author	Devika Sharma and Zvika Brakerski (WEI)

Abstract. This document discusses the results achieved by the partners towards constructing and defining building blocks for practical advanced protocols, in the half-life of PROMETHEUS WP4. More specifically, it gives a detailed technical overview of the results in the 11 research papers covered here and provides a summary of relevant open questions.

Keywords: Lattice based privacy-preserving cryptography, building blocks..

Signatures

Written by	Devika Sharma and Zvika Brakerski	WEI	December 2019
Reviewed by	Leo Ducas	CWI	17/122019
Reviewed by	Javier Herranz	UPC	18/12/2019
Approved by	Benoît Libert as Project coordinator	ENSL	23/12/2019
Approved by	Sébastien Canard as Technical leader	ORA	23/12/2019

Partners

ENSL	ENS de Lyon
ORA	Orange SA
CWI	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	Nederlandse organisatie voor Toegepast-Natuurwetenschappelijk Onderzoek
UPC	Universitat Politècnica de Catalunya · BarcelonaTech
UR1	Université de Rennes 1
WEI	Weizmann Institute of Science

Contents

1	Introduction	6
1.1	Progress towards objectives	6
2	Signatures	10
2.1	Masking the GLP Lattice-Based Signature Scheme [BBE ⁺ 18]	10
3	Pseudorandom Functions	12
3.1	Adaptively Secure Distributed PRFs from LWE [LST18]	12
4	Functional Encryption	14
4.1	Multi-Client Functional Encryption from LWE [LT19]	14
4.2	Fully Secure ABE for t-CNF from LWE [Tsa19]	18
5	Homomorphic Encryption	20
5.1	Rate-1 Fully-Homomorphic Encryption [BDGM19]	20
6	Zero-knowledge	23
6.1	Lattice-Based Zero-Knowledge Arguments [LLNW18]	23
6.2	Zero-Knowledge Elementary Databases [LNTW19]	24
6.3	RLWE-based Zero-Knowledge Proofs [MM19]	26
7	Implementation	28
7.1	Implementing RLWE-based Schemes [AHH ⁺ 19]	28
7.2	A Comparison of the Homomorphic Encryption Libraries [MKLR18]	30
7.3	On Standardising Sparse-secret LWE Parameter Sets for Homomorphic Encryption [CP19]	33
8	Conclusion	36

List of Tables

1	Communication cost (in bits)	27
---	--	----

List of Figures

1	Commitment's size of Xie <i>et al.</i> (\ominus), Benhamouda <i>et al.</i> (\boxplus) and Martínez <i>et al.</i> (\triangleleft)	27
2	Evolution of the noise as a function of the multiplying depth when $\log p = 64$. For $\log p = 256$ and $\log p = 2048$, results are similar. Results for HELib-MP are given twice, with and without the special primes used in the relinearization operation.	31
3	Average time for one multiplication as a function of the multiplying depth when $\log p = 256$	32
4	Average time for one multiplication as a function of the multiplying depth when $\log p = 2048$	32
5	An estimate of the ring operations (rop) required to solve the LWE instances parameterised by $n = 1024$, $q = 2^{40}$ and $\sigma \approx 3.2$ and a sparse ternary secret with hamming weight $h \in \{64, 128, 256, 512\}$, using the <i>usvp</i> , <i>dual</i> , <i>hybrid-dual</i> and <i>hybrid-dec</i> attacks.	34
6	Required bitsize $\log q$ to achieve a security level of λ for an LWE instance parameterised by dimension n , modulus q , error standard deviation $\sigma = 3.19$ and a uniform ternary secret with hamming weight $h = \lambda$, under the lattice reduction cost model $T_{\text{BKZ}}(\beta, d) = 2^{0.292\beta + 16.4 + \log(8d)}$. The solid lines represent data points and the dashed lines represent extrapolation to include $n = 65536$ and $n = 131072$	35

1 Introduction

Work package 4 (WP4) of the PROMETHEUS project pursues two main objectives: extending the reaches of practical lattice-based cryptography and exploring novel constructions with functionalities that are not currently present in any (not necessarily lattice-based) primitive. As is well-known, there are very efficient basic (lattice-based) primitives such as public-key encryption, identity-based encryption and digital signatures, but not as many efficient constructions for practical, more advanced primitives. This work package aims to provide better lattice-based signatures, pseudorandom functions compatible with zero-knowledge proofs, and more efficient realisations of lattice-based homomorphic commitment schemes. It, further, plans to improve the efficiency of fully homomorphic encryption with post-quantum security guarantees and also improve the solutions (in terms of efficiency, expressiveness, and security) of post-quantum access-control mechanisms such as threshold encryption and attribute-based encryption.

The results in this work package do not merely try to translate the number-theoretic constructions into lattice-based ones, as this approach may not result in efficient constructions. In fact, the attempt is to illuminate which types of constructions can be efficiently built from lattices, and then try to design efficient building blocks that would facilitate constructing practical and advanced protocols in work packages 5 and 6. With this understanding in mind, the goal is to have novel protocol designs that are based on the types of basic building blocks that can be built efficiently from lattices, by the end of the project. The January 2018 issue of the ERCIM NEWS magazine, specially dedicated to the theme of ‘Quantum Computing’, published an article ([VAvHD18]) describing the threats the advent of quantum computers would pose on the current communication and how PROMETHEUS aims to be ready with solutions. Quote “In the post-quantum era, most of the currently used cryptography is no longer secure due to quantum attacks. Cryptographers are working on several new branches of cryptography that are expected to remain secure in the presence of a universal quantum computer. Lattice-based cryptography is currently the most promising of these branches. The new European PROMETHEUS project will develop the most secure design and implementations of lattice-based cryptographic systems. Exploitation of the project results will be stimulated by demonstrating and validating the techniques in industry-relevant environments.” Unquote. This deliverable is evidence that PROMETHEUS is on the right track.

1.1 Progress towards objectives

Regarding the scientific achievements of WP4, the partners have pursued three distinct research directions in WP4. The first one relates to the design and the provable security of lattice-based primitives that can serve as building blocks (Task 4.1 for signatures and Task 4.2 for encryption) for higher-level privacy-preserving protocols (in WP5 and 6). The second one focuses on zero-knowledge proofs (Task 4.3) allowing to prove statements in lattice-related languages while preserving the secrecy of provers’ inputs (which is in particular necessary in most of privacy-preserving protocols). The last one is about the implementation of quantum-safe cryptographic primitives and their security (Task 4.4, which is a necessary input to WP6).

1.1.1 Task 4.1 Lattice based signatures and other building blocks

The design of privacy-preserving protocols usually requires digital signatures that are compatible with zero-knowledge proofs. Efficient lattice-based signature schemes exist, but only in the random oracle model. Such schemes are unfortunately not well-suited for our applications because zero-knowledge proofs do not smoothly interact with cryptographic hash functions such as SHA-256. The main reason for the latter is that it is rather inefficient to prove the knowledge of a pre-image of a hash function that does not have some algebraic structure. We thus need signature schemes with security proofs in the standard model. Moreover, applications like anonymous credentials or e-cash systems require a signature flavour called “signatures with efficient protocols” which supports efficient two-party protocols, as explained in D4.1.

An important part of Task 4.1 is to come up with better realisations of secure lattice-based signatures with efficient protocols. Towards this end, work done in [BBE⁺18] describes a modified but much more secure construction of the existing GLP signature scheme [GLP12]. More elaborately, the authors in this work show how to efficiently mask the (key generation and signing algorithm in the) GLP scheme at any order so as to achieve security against power analysis and related attacks (both simple power analysis and higher-order attacks like differential/correlation power analysis). This work is the first instance of masking being applied to a lattice-based signature scheme. The masked signature scheme is EUF-CMA (existential unforgeability under chosen message attack) secure in the threshold probing model (ISW, [ISW03]), wherein the adversary can read off at most d wires in a circuit. This model is equivalent [DDF14] to the more realistic noisy model where the adversary acquires leakage on *all* variables, but that leakage is perturbed with some noise, as in the case of practical side-channel attacks.

As building blocks for the design of e-cash solutions in WP5, WP4 aims to build better lattice-based pseudo-random functions (PRFs) that can be smoothly combined with zero-knowledge proofs (see PROMETHEUS’ deliverable D4.1 but also D5.2). In particular, a prover should be able to convince a verifier that some value is the correct evaluation of a PRF for some committed (or encrypted) inputs and keys. While such statements can be efficiently handled under discrete logarithm assumptions, no quantum-resistant solution were known so far. In the first of its kind, [LST18] constructs a non-interactive *adaptively secure* distributed PRF in the standard model. This construction is secure under the LWE assumption with super-polynomial approximation factors against adversaries that may adaptively decide which servers to corrupt. Also, as a by-product of work done in [Tsa19], we see a construction of a lattice-based, fully secure single-key constrained PRF from OWF for a particular function class.

1.1.2 Task 4.2 Lattice-based encryption schemes with additional properties

This task aims to address the functionality advantages offered by lattice-based cryptography in the context of encryption schemes with advanced properties, such as the feasibility of computing over encrypted datasets. In addition, this task considers the extent to which certain existing public-key functionalities can be adapted to the world of lattices. Addressing privacy issues via encryption raises major challenges if we want to maintain the ability to process encrypted data. Modern study shows that lattices are promising tools for this purpose, as they enable cryptocomputing functionalities such as fully homomorphic and functional encryption. Lattices thus provide the double benefit of increased security and enhanced functionality. In the setting

of encryption schemes with advanced properties, [LT19] constructs a Multi-Client Functional Encryption (MCFE) schemes for linear functions. A MCFE scheme supports the evaluation of multivariate functions over data coming from distinct sources. This work gives the first construction of a standard-model MCFE scheme that is fully secure in an adaptive corruption setting under the well-studied LWE assumption. It also provides a decentralized variant of this scheme and shows that it is secure in the static corruption setting, but for adaptively chosen messages. Both constructions are proved secure under the LWE assumption with sub-exponential approximation factors. Functional encryption is one possible primitive for PROMETHEUS' Cyber Threat Intelligence use case.

Further in this direction, [Tsa19] considers ciphertext-policy attribute-based encryption (CP-ABE), where ciphertexts are labeled with an access policy and can only be decrypted by keys associated with attributes that satisfy the access policy of the ciphertext. This work provides for the first time a lattice-based (ciphertext-policy) ABE scheme for the function class t -CNF, which consists of CNF formulas where each clause depends on at most t bits of the input, for any constant t . This class includes NP-verification policies, bit-fixing policies and t -threshold policies. Even if this result is not directly related to a PROMETHEUS use case, ABE has enough potential to be a candidate for standardization (especially at the ETSI level, as closely related to Identity Based Encryption).

One of the focal points of task 4.2 is better solutions for fully homomorphic encryption (FHE), one of the triumphs of lattice-based cryptography, which turned from fantasy into reality in under 10 years. FHE allows performing computation on encrypted data without decrypting it first, and is thus one of the basic desired tasks in a world where computation is performed remotely. In addition, FHE is known to imply short non-interactive zero-knowledge proofs for any NP statement (assuming the existence of NIZK for simple statements): namely, the size of the proof only depends on the length of the witness. As such, more efficient FHE realisations are likely to positively impact the protocols to be developed in WP5. Toward the objective of achieving more efficient FHE schemes, [BDGM19] constructs an optimal-rate (rate-1) FHE scheme that is secure under the LWE assumption with polynomial modulus-to-noise ratio. FHE is a primitive related to functional encryption and multi-party computation, and then an important building block to the Cyber Threat Intelligence use case.

1.1.3 Task 4.3 Lattice-based zero-knowledge proofs of knowledge

Zero-knowledge proofs are at the heart of every privacy-preserving protocol, as explained in several other PROMETHEUS' deliverables, such as D4.1, D5.1 and D5.2. Some examples of the types of things that often need to be proved are: the knowledge that the public key is validly constructed; the knowledge of the plaintext encrypted in the ciphertext (e.g., in anonymous credentials); knowledge of a signature of a message (e.g., in anonymous credentials, e-cash and e-voting); proving that the plaintext is in a certain range (e.g., in anonymous credentials, e-cash and e-voting); and proving that mix-net shuffling was correctly done (e.g., in e-voting). Such proof constructions have been well studied for classical protocols and are very efficient under assumptions like RSA or Diffie-Hellman. While they have analogous realisations based on the hardness of LWE and Ring-LWE problems, these are not so efficient.

Work done in [LLNW18], constructs Zero-Knowledge arguments to prove integer-relations among commitments. More precisely, this work gives statistical zero-knowledge

arguments allowing a prover to convince a verifier that x , y and z are commitments to integers X , Y and Z , respectively that satisfy additive, multiplicative, order and range relations. These arguments are secure under the standard (non-ideal) LWE assumption with both polynomial moduli and approximation factors.

Further, work done in [LNTW19], constructs zero-knowledge proof techniques allowing to protect the privacy of sensitive databases. In zero-knowledge databases, a prover commits to an elementary database (i.e., a set of key-value pairs (x, y) where each key x has at most one value y) and subsequently proves statements about the committed database without even revealing the database size. Previously, all non-interactive such protocols were limited to proving simple statements such as the membership or the non-membership of specific elements x . This work describes techniques that allow a prover to prove more general statements, including range queries (i.e., provably reveal all database keys in a specific range $[a, b]$). These arguments are secure under standard lattice assumptions.

Finally, work done in [MM19] investigates the design of efficient Zero-Knowledge Proofs of Knowledge for linear and multiplicative relations among messages committed using a Ring Learning With Errors (RLWE) based commitment scheme. This 5-move protocol achieves perfect zero-knowledge, reduces the communication cost from previous Stern-based schemes, but still incurs a soundness error of approximately $1/2$. Reducing the soundness error to a negligible upper bound thus requires parallel repetitions which seriously hurt the efficiency.

1.1.4 Task 4.4 Implementation of building block

Besides suitable cryptographic properties another important aspect is the practicality of the investigated lattice-based schemes. This needs to be evaluated by implementation and by evaluating different security levels for a range of different target platforms. Primarily, this includes the reference implementation of the identified lattice-based schemes (WP4.1 to WP4.3) in software, to be operated and evaluated on common processor platforms. In future applications for the Internet of Things (IoT), however, a majority of devices will still be based on significantly smaller processors that are often severely constrained in their features, including processing power, memory or energy consumption. Based on the provided reference implementations, a secondary goal of this WP is to evaluate the identified schemes for such constraints for low-cost embedded software devices. To this end, [AHH⁺19], implemented several public-key encryption schemes based on the “Ring Learning-With-Errors” (RLWE) assumption using an RSA co-processor. They notably report an implementation of a module-LWE based key encapsulation scheme on a smart card which is equipped with an RSA co-processor. The results demonstrate comparable performance to running RSA itself. In a similar flavour, authors of [MKLR18] provide a comparative benchmark of the leading homomorphic encryption libraries HELib, FV- NTLlib, and SEAL for large plaintext moduli of up to 2048 bits, and analyse their relative performance. Further exploring secure parameters for implementing Homomorphic encryption, [CP19] discusses the security of possible sparse-secret LWE parameter sets against hybrid attacks. The authors in this work present a conservative analysis of the hybrid attacks for parameter sets of varying sparsity, with the goal of balancing security requirements with bootstrapping efficiency. They also argue that the methodology in the Homomorphic Encryption Security Standard, as published by the HomomorphicEncryption.org consortium, can be easily modified to support dimensions higher than the current (fixed) upper bound.

Structure: In what follows, we broadly divide the work done under PROMETHEUS until now into six sections; Signatures, Psuedorandom Functions, Functional Encryption, Homomorphic Encryption, Zero-Knowledge Proofs and Implementation. Each section contains a brief technical overview of the relevant paper(s) and concludes with a short discussion on open questions.

2 Signatures

Lattice-based cryptography is an attractive option in the post-quantum setting, as it allows designing post-quantum implementations of a wide range of primitives with strong security guarantees and a level of efficiency comparable to currently deployed RSA and elliptic curve-based schemes. However, it poses new sets of challenges as far as side-channels and other physical attacks are concerned. For instance, as demonstrated in [BHLY16], a cache attack targeting the Gaussian sampling of the randomness used in BLISS signatures can recover the entire secret key from the side-channel leakage of a few thousand signature generations. This makes BLISS signatures unfavourable for implementation even though their performance and key and signature sizes are comparable to RSA and ECDSA signatures, as claimed in [DDLL13]. As the call for concrete implementations and deployment of postquantum cryptography becomes more pressing, safe guarding against such attacks is a crucial issue to address. To this end, a few countermeasures have been proposed. In particular, the powerful technique of masking has been applied to the decryption procedure of some lattice-based encryption schemes. Masking, a well-known technique introduced in [CRR02], essentially consists of splitting a secret value into $d + 1$ values (d is the masking order), using a secret sharing scheme. This forces an adversary to read many internal variables if he wants to recover the secret value, and he will gain no information if he observes fewer than d values.

However, it is not always easy to mask an implementation of a cryptographic scheme. Some difficulties specific to attempting to mask the BLISS signatures, for instance, are discussed in the paper [BBE⁺18]. Largely, it is the non-linearity and the Gaussian randomness involved in most signature schemes that hinder masking their protocols efficiently. Although, there exist lattice-based signatures, for eg. the GLP scheme [GLP12], that entirely avoid Gaussians and other distributions, and hence seem to support masking in a more natural way.

2.1 Masking the GLP Lattice-Based Signature Scheme [BBE⁺18]

In this paper, the authors show how to efficiently mask the (key generation and signing algorithm in the) GLP scheme at any order so as to achieve security against power analysis and related attacks (both simple power analysis and higher-order attacks like differential/correlation power analysis). This work is the first instance of masking being applied to a lattice-based signature scheme. The masked signature scheme is EUF-CMA (existential unforgeability under chosen message attack) secure in the threshold probing model (ISW, [ISW03]), wherein the adversary can read off at most d wires in a circuit. This model is equivalent [DDF14] to the more realistic noisy model where the adversary acquires leakage on *all* variables, but that leakage is perturbed with some noise, as in the case of practical side-channel attacks.

We briefly describe the original GLP protocol and its masked version. For the setup, let n be a power of 2, p a prime congruent to 1 modulo $2n$ and the ring $\mathcal{R} := \mathbb{Z}_p[x]/\langle x^n + 1 \rangle$. The elements of \mathcal{R} can be represented by polynomials of degree $n - 1$ with coefficients in the range $[-(p-1)/2, (p-1)/2]$. For an integer k such that $0 < k \leq (p-1)/2$, we denote by \mathcal{R}_k the elements of \mathcal{R} with coefficients in the range $[-k, k]$. Let $H : \{0, 1\} \rightarrow \mathcal{D}_\alpha^n$ be a particular cryptographic hash function. Here \mathcal{D}_α^n is the set of polynomials in \mathcal{R} that have all zero coefficients except for at most $\alpha = 32$ coefficients that are in $\{-1, 1\}$. The parameter k controls the trade-off between the security and the run time of the scheme. The smaller k gets, the more secure the scheme becomes along with shorter signatures but the time to sign increases. For a scheme of masking order d ,

Key Generation: Outputs Signing key sk and verification key pk as follows;

1. Generate secret keys s_1 and s_2 in \mathcal{R}_1 in their masked form $(s_{1,i})_{0 \leq i \leq d}$ and $(s_{2,i})_{0 \leq i \leq d}$, respectively.
2. Choose $\mathbf{a} \xleftarrow{\$} \mathcal{R}$
3. For $0 \leq i \leq d$, compute $t_i = as_{1,i} + s_{2,i}$ and $t = \sum_{i=0}^d t_i$
4. Return $sk = ((s_{1,i})_i, (s_{2,i})_i)$ and $pk = (\mathbf{a}, t)$.

Signing: Given message m , pk and sk , output signature σ as follows;

1. Generate y_1 and y_2 in \mathcal{R}_k in their masked form $(y_{1,i})_{0 \leq i \leq q}$ and $(y_{2,i})_{0 \leq i \leq q}$
2. For each $0 \leq i \leq d$, compute $r_i = ay_{1,i} + y_{2,i}$ and $r = \sum_i r_i$.
3. Compute $c = H(r, m)$ and $z_{1,i} = cs_{1,i} + y_{1,i}$ and $z_{2,i} = cs_{2,i} + y_{2,i}$, for $0 \leq i \leq q$.
4. Run the masked version of rejection sampling to check if $z_1 = \sum_i z_{1,i}$ and $z_2 = \sum_i z_{2,i}$ lie in $\mathcal{R}_{k-\alpha}$. See BBE⁺18, Algorithm 16 for details.
5. Return $\sigma = (z_1, z_2, m)$.

Verification: Given m , σ , pk , accept if $z_1, z_2 \in \mathcal{R}_{k-\alpha}$ and $c = H(az_1 + z_2 - tc, m)$, otherwise reject.

The above protocol reveals the value (r, c) , even if the execution is rejected. Intuitively, this would pose a threat as a side-channel attacker can obtain information about the secret from the values corresponding to the rejected samples. However, this is common practice in Zero-Knowledge proofs, where the prover sends the commitment r , then the verifier samples and sends a challenge c and the prover finally computes the response z . Consequently, these values are public in authentication schemes. In the case of Fiat-Shamir with Aborts, Vadim Lyubashevsky [Lyu09] has proved that the output z is independent of the secret even though the values (r, c) are revealed. The other solution is to use a commitment scheme to hide these values before computing the response z . This, however, makes the scheme more complicated to mask.

The masked signature scheme described above is EUF-CMA (existential unforgeability under chosen message attack) secure in the threshold probing model (ISW, [ISW03]), wherein the adversary can read off at most d wires in a circuit. This model is equivalent [DDF14] to the more realistic noisy model where the adversary acquires

leakage on *all* variables, but that leakage is perturbed with some noise, as in the case of practical side-channel attacks.

The authors implemented their results on an Intel Core CPU and it was observed that, for order $d = 1, 2, 3$, the overhead in running time of the masked scheme was around $15\times$, $30\times$ and $73\times$, respectively. Even though the parameters chosen in the paper are not optimized, these results are promising. The paper includes several suggestions that could speed up implementation.

Open Problems: It would be a useful attempt to try and apply this method to other lattice-based Fiat-Shamir type signature schemes that use uniform distributions in intervals (as opposed to Gaussian distributions). On the other hand, developing such a framework for schemes that involve Gaussian distributions still remains a formidable challenge. Finally, it would be interesting to leverage recent advances in verification and synthesis of masked implementations in a more systematic way in the lattice-based setting. For instance, the sheer size of the verification algorithms involved poses significant challenges in terms of scalability; however, automated tool support would be invaluable for the further development of masking in the postquantum setting.

3 Pseudorandom Functions

A pseudorandom function (PRF) family is a set F of keyed functions with common domain Dom and range Rng such that no PPT adversary can distinguish a real experiment, where it has oracle access to a random member $f \leftarrow F$ of the PRF family, from an ideal experiment where it is interacting with a truly random function $R : \text{Dom} \rightarrow \text{Rng}$. To be useful, a PRF should be efficiently computable - meaning that $F_s(x)$ must be deterministically computable in polynomial time given the key s and the input $x \in \text{Dom}$ - and the key size must be polynomial. The following work deals with a variant of it, called distributed PRF.

3.1 Adaptively Secure Distributed PRFs from LWE [LST18]

In a (threshold) distributed PRF (DPRF), secret keys are broken into N shares s_1, \dots, s_N , each of which is given to a different server. Using its secret key share s_i , the i -th server locally computes a partial evaluation $F_{s_i}(x)$ of the function. A dedicated server then gathers at least $t \leq N$ correct partial evaluations $F_{s_{i_1}}(x), \dots, F_{s_{i_t}}(x)$ and reconstructs the evaluation $F_s(x)$ for the long-term key s . DPRFs inherit the usual benefits of threshold cryptography; when $t < N$, firstly, they allow fault-tolerant systems to keep running even when some of the servers crash, and secondly, the adversary is forced to break into t servers to compromise the security of the whole scheme. Here, we restrict our discussion to non-interactive constructions whose security in the standard model is based on lattice assumptions. All such constructions known thus far are secure in the static corruption setting, where the adversary chooses the servers to corrupt at the beginning of the game, before any evaluation query.

For a polynomial N and when $t \approx N/2$, proving adaptive security is considerably more challenging as a trivial complexity leveraging argument (i.e., guessing the set of corrupted servers upfront) makes the reduction super-polynomial. Moreover, as shown in [LST18], allowing a *single* partial evaluation query before the first corruption query already results in a definition which is strictly stronger than that of static security. In the adaptive corruption setting, the difficulty is that, by making N partial

evaluation queries before corrupting any server, the adversary basically commits the challenger to all secret key shares. Hence, a reduction that only knows $t - 1 \approx N/2$ shares is unlikely to work as it would have to make up its mind on which set of $t - 1$ shares it wants to know at the outset of the game. In particular, this hinders a generic reduction from the security of an underlying key-homomorphic PRF. This suggests to find a reduction that knows all shares of the secret key, making it easier to consistently answer adaptive corruption queries.

To this end, the partners [LST18] turn to lossy trapdoor functions [PW08], which are function families that contain both injective and lossy functions with computationally indistinguishable evaluation keys. The construction of [LST18] relies on the fact that the LWE function and its deterministic LWR variant [BPR12] are both lossy trapdoor functions (as shown in [GKPV10, BKPW12, AKPW13]). Namely, the function that maps $\mathbf{s} \in \mathbb{Z}^n$ to¹ $\lfloor \mathbf{A} \cdot \mathbf{s} \rfloor_p$ is injective when $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ is a random matrix and becomes lossy when \mathbf{A} is of the form $\bar{\mathbf{A}} \cdot \mathbf{C} + \mathbf{E}$, where $\bar{\mathbf{A}} \in \mathbb{Z}_q^{m \times n'}$, $\mathbf{C} \in \mathbb{Z}_q^{n' \times n}$ are uniformly random and $\mathbf{E} \in \mathbb{Z}^{m \times n}$ is a small-norm matrix. The idea of Libert, Stehlé and Titu [LST18] is to first construct a PRF which maps an input x to $\lfloor \mathbf{A}(x) \cdot \mathbf{s} \rfloor_p$, where $\mathbf{s} \in \mathbb{Z}^n$ is the secret key and $\mathbf{A}(x) \in \mathbb{Z}_q^{m \times n}$ is derived from public matrices. The construction of [LST18] thus evaluates a lossy trapdoor function on an input consisting of the secret key using a matrix that depends on the input. In the security proof, [LST18] uses admissible hash functions [BB04] and techniques from fully homomorphic encryption [GSW13a] to “program” $\mathbf{A}(x)$ in such a way that, with non-negligible probability, it induces a lossy function in all evaluation queries and an injective function in the challenge phase.² (This use of lossy trapdoor functions is somewhat unusual since their injective mode is usually used to handle adversarial queries while the lossy mode comes into play in the challenge phase.) By choosing a large enough ratio q/p , [LST18] can make sure that evaluation queries always reveal the same information about the secret \mathbf{s} . Since $\lfloor \mathbf{A}(x^*) \cdot \mathbf{s} \rfloor_p$ is an injective function in the challenge phase, the security proof of [LST18] argues that the secret key has high min-entropy, even conditionally on responses to evaluation queries. At this point, they can extract statistically uniform bits from $\lfloor \mathbf{A}(x^*) \cdot \mathbf{s} \rfloor_p$ using a deterministic randomness extractor: analogously to the deterministic encryption case [RSV13], the proof of [LST18] needs to handle a source that may be correlated with the seed.

The above approach bears resemblance with key-homomorphic PRFs [BLMR13, BP14] which also evaluate functions of the form $\lfloor \mathbf{A}(x) \cdot \mathbf{s} \rfloor_p$. However, our proof method is very different in that it relies on the lossy mode of LWE and the homomorphic encryption scheme of [GSW13a]. The advantage of the approach taken in [LST18] is that the challenger knows the secret key \mathbf{s} at all steps of the security proof. In the distributed setting, this makes it easier to handle adaptive adversaries because the reduction can always correctly answer corruption queries. In order to share the secret key \mathbf{s} among N servers, the partners [LST18] rely on the Linear Integer Secret Sharing (LISS) schemes of Damgård and Thorbek [DT06], which nicely fit the requirements of their security proof. Among other properties, they allow secret key shares to remain small with respect to the modulus, which helps [LST18] make sure that partial evaluations – as lossy functions of their share – always reveal the same information about uncorrupted shares. Moreover, LISS also enable small reconstruction constants: the secret \mathbf{s} can be recovered as a linear combination of authorized shares

¹Introduced in [BPR12], the notation $\lfloor x \rfloor_p$ stands for the rounded value $\lfloor (p/q) \cdot x \rfloor \in \mathbb{Z}_p$, where $x \in \mathbb{Z}_q$, and $p < q$.

²They use a “find-then-guess” security game where the adversary obtains correct evaluation for inputs of its choice before trying to distinguish a real function evaluation from a random element of the range.

with coefficients in $\{-1, 0, 1\}$, which is useful to avoid blowing up error terms when partial evaluations are combined together. A notable difference with [DT06] is that the DPRF of [LST18] uses a LISS scheme with Gaussian entries (instead of uniform ones), which makes it easier to analyze the remaining entropy of the key in the final step of the security proof.

Open Problems: It remains an open problem to define threshold PRFs under standard lattice assumptions with polynomial approximation factors. Another challenging open problem that would be very relevant to the current flavour of research is to find pseudorandom functions that can smoothly interact with more efficient zero-knowledge protocols based on the ‘Fiat-Shamir with aborts’ technique. Currently, the latter would not enable a knowledge extractor that really extracts witnesses containing the PRF seed and its input.

4 Functional Encryption

Functional encryption (FE) is a modern paradigm that overcomes the all or nothing nature of ordinary encryption schemes. In FE, the master secret key \mathbf{msk} allows deriving a sub-key \mathbf{dk}_f associated with a specific function f . When \mathbf{dk}_f is used to decrypt a ciphertext C of a message X , the decryptor only obtains $f(X)$ and learns nothing else about X . In the case of FE for linear functions, the constructions described in [ABCP15, ALS16] are secure under the LWE assumption, against adaptive adversaries. Functional encryption is an extremely general concept as it subsumes identity-based encryption, searchable encryption, attribute-based encryption, broadcast encryption and many others. We, particularly, discuss two variants of it here.

Many natural applications of FE require computing over data coming from multiple parties. In such a setup, ideally, the participants should be able to supply their input without interacting with one another and go off-line immediately after contributing their share. This motivates the concept of multi-client functional encryption (MCFE) as described in [GGJS13, GKL⁺13]. Section 4.1 describes a MCFE construction.

Section 4.2 describes an ABE construction. Attribute-based Encryption (ABE), is a public key encryption system that can support multiple users with varying decryption permissions. In this work [Tsa19], the authors focus on ciphertext-policy ABE schemes, where each ciphertext is associated with a public policy f and each decryption key is associated with a public attribute x , such that decryption succeeds conditioned on $f(x) = 1$.

4.1 Multi-Client Functional Encryption from LWE [LT19]

MCFE supports the evaluation of multivariate functions over data coming from distinct sources. More precisely, it allows ℓ clients to encrypt ciphertexts $(\mathbf{C}_{t,1}, \dots, \mathbf{C}_{t,\ell})$ under some label (or tag) t . These tags may correspond to time specific information or a dataset name. Each client can encrypt his own data X_i for a label t using a private encryption key s_i issued by a trusted authority in such a way that, as long as all $\mathbf{C}_{t,i}$ share the same label t , an evaluator endowed with a functional key \mathbf{dk}_f can evaluate $f(X_1, \dots, X_\ell)$ without learning anything else on the underlying plaintexts X_i . Functional decryption keys can be derived by the central authority using the master secret key.

In certain scenarios, where the clients may be reluctant to rely on a single party of trust, a decentralized version of MCFE is used. Decentralized multi-client functional encryption (DMCFE) obviates the need for a centralized authority by shifting the task of generating functional secret keys to the clients themselves. In the setup phase, the clients $\mathcal{S}_1, \dots, \mathcal{S}_\ell$ first generate public parameters by running an interactive protocol but no further interaction is needed among clients when it comes to generating functional secret keys later on. When a decryptor wishes to obtain a functional secret key for an ℓ -ary function f , it interacts with each client \mathcal{S}_i independently so as to obtain partial functional decryption keys $\mathbf{dk}_{f,i}$. The decryptor can then fold $\{\mathbf{dk}_{f,i}\}_{i=1}^\ell$ into a functional decryption key \mathbf{dk}_f for f . In this scenario, each client has full control over his individual data and the functions for which secret keys are given out. Furthermore, no interaction among senders is required beyond the setup phase, where public parameters are generated.

A recent work, [CDG⁺18], describes an adaptively secure (D)MCFE scheme for evaluating linear functions over integers. This construction is adaptively secure in the random oracle model under the Decisional Diffie-Hellman (DDH) assumption. In the standard model, the original MCFE construction [GKL⁺14] is only known to be statically secure and relies on indistinguishability obfuscation.

In [LT19], the authors give the first construction of a standard-model MCFE scheme that is fully secure in an adaptive corruption setting under the well-studied LWE assumption. They also provide a decentralized variant of their scheme and show that it is secure in the static corruption setting, but for adaptively chosen messages. Both constructions are proved secure under the LWE assumption with sub-exponential approximation factors. This construction is inspired by but not an analogue of the MCFE scheme of [CDG⁺18].

As defined in [SGG]⁺14, GKL⁺14], multi-client functional encryption allows computing over input vectors (X_1, \dots, X_ℓ) where coordinate X_i may be sent by a different client. Each ciphertext C_i is associated with a client index i and a tag t (also called “label”). On input of ciphertexts $(C_1 = \text{Encrypt}(1, X_1, t), \dots, C_\ell = \text{Encrypt}(\ell, X_\ell, t))$, where C_i is generated by client i using a secret encryption key ek_i for each $i \in [\ell]$, anyone holding a functional decryption key dk_f for an ℓ -ary function can compute $f(X_1, \dots, X_\ell)$ as long as all C_i are labeled with the same tag t (which may be a time-specific information or a dataset name). No further information than $f(X_1, \dots, X_\ell)$ is revealed about individual inputs X_i and nothing can be inferred by combining ciphertexts generated for different tags.

The construction of [LT19] computes linear combinations of vectors encrypted by $(C_1 = \text{Encrypt}(1, X_1, t), \dots, C_\ell = \text{Encrypt}(\ell, X_\ell, t))$. It starts from the observation that the DDH-based MCFE scheme of Chotard *et al.* [CSG⁺18] can be interpreted as relying on (a variant of) the key-homomorphic pseudorandom function [BLMR13] of Naor, Pinkas and Reingold [NPR99]. Namely, the scheme of [CSG⁺18] encrypts $x_i \in \mathbb{Z}_q$ for the tag t by computing $C_i = g^{x_i} \cdot H_{t,1}^{s_i} \cdot H_{t,2}^{t_i}$, where $(s_i, t_i) \in \mathbb{Z}_q^2$ is the i -th sender’s secret key and $(H_{t,1}, H_{t,2}) = H(t) \in \mathbb{G}^2$ is derived from a random oracle in a DDH-hard group $\mathbb{G} = \langle g \rangle$.

The security proof of [CSG⁺18] crucially exploits the entropy of the secret key (s_i, t_i) in a hybrid argument over all encryption queries. To preserve this entropy, they need to prevent the encryption oracle from leaking too much about uncorrupted users’ secret keys $\{(s_i, t_i)\}_i$. For this purpose, they rely on the DDH assumption to modify the random oracle $H : \{0, 1\}^* \rightarrow \mathbb{G}^2$ in such a way that, in all encryption queries but one, the hash value $H(t) \in \mathbb{G}^2$ lives in a one-dimensional subspace.

A natural idea is to replace the random-oracle-based key-homomorphic PRF of

[NPR99] by an LWE-based key-homomorphic PRF [BLMR13, BP14]. However, analogously to Chotard *et al.* [CSG⁺18],³ the partners aim at an MCFE system that can be proved secure in a game where the adversary is allowed to corrupt senders adaptively. In order to deal with the adaptive corruption of senders, the scheme of [LT19] thus relies on the adaptively secure distributed PRF proposed by Libert, Stehlé and Titu [LST18]. The latter can be seen as instantiating the programmable hash function [HK08] of Freire *et al.* [FHPS13] in the context of homomorphic encryption (FHE). Their PRF maps an input x to $\lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p$, where $\mathbf{s} \in \mathbb{Z}^n$ is the secret key and $\mathbf{A}(x) \in \mathbb{Z}_q^{n \times m}$ is derived from public matrices using the Gentry-Sahai-Waters FHE [GSW13a]. More precisely, the matrix $\mathbf{A}(x)$ is obtained as the product of GSW ciphertexts dictated by the output of an admissible hash function [BB04] applied to the PRF input. The security proof of the distributed PRF in [LST18] uses the property that, with noticeable probability, the input-dependent matrix $\mathbf{A}(x)$ is a GSW encryption of 1 for the challenge input x^* : namely, $\mathbf{A}(x^*)$ is a matrix of the form $\mathbf{A}(x^*) = \mathbf{A} \cdot \mathbf{R}^* + \mathbf{G}$, where $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$ is the gadget matrix of Micciancio and Peikert and $\mathbf{R}^* \in \mathbb{Z}^{m \times m}$ is a small-norm matrix. At the same time, all evaluation queries are associated with a matrix $\mathbf{A}(x)$ consisting of a GSW encryption of 0 (i.e., a matrix $\mathbf{A}(x) = \mathbf{A} \cdot \mathbf{R}$, for a small-norm $\mathbf{R} \in \mathbb{Z}^{m \times m}$). Then, the proof of [LST18] appeals to the lossy mode of LWE [GKPV10] and replaces the uniform matrix $\mathbf{A}^\top \in \mathbb{Z}_q^{m \times n}$ by a lossy matrix of the form $\hat{\mathbf{A}}^\top \cdot \mathbf{C} + \mathbf{E}$, where $\mathbf{E} \in \mathbb{Z}^{m \times n}$ is a short integer matrix with Gaussian entries, $\mathbf{C} \in \mathbb{Z}_q^{n_1 \times n}$ is random, and $\hat{\mathbf{A}} \in \mathbb{Z}_q^{n_1 \times m}$ has rank $n_1 \ll n$. In all evaluation queries, the smallness of $\mathbf{s} \in \mathbb{Z}^n$ then ensures that the values $\lfloor \mathbf{A}(x)^\top \cdot \mathbf{s} \rfloor_p$ always reveal the same information about \mathbf{s} , which amounts to the product $\mathbf{C} \cdot \mathbf{s} \in \mathbb{Z}_q^{n_1}$. Since $\mathbf{A}(x^*)$ depends on \mathbf{G} for the challenge input x^* , the function $\lfloor \mathbf{A}(x^*)^\top \cdot \mathbf{s} \rfloor_p$ is in fact an injective function of \mathbf{s} , meaning that it has high min-entropy.

The MCFE scheme of [LT19] relies on the lossy mode of LWE in a similar way to [LST18], except that it adds a Gaussian noise instead of using the Learning-With-Rounding technique [BPR12]. The i -th sender uses his secret key $\mathbf{s}_i \in \mathbb{Z}^n$ to encrypt a short integer vector as $\vec{x}_i \in \mathbb{Z}^{n_0}$ as $\mathbf{C}_i = \mathbf{G}_0^\top \cdot \vec{x}_i + \mathbf{A}(t)^\top \cdot \mathbf{s}_i + \text{noise} \in \mathbb{Z}_q^m$, where $\mathbf{A}(t) \in \mathbb{Z}_q^{n \times m}$ is a tag-dependent matrix derived as a product of GSW ciphertexts indexed by the bits of t and $\mathbf{G}_0 \in \mathbb{Z}_q^{n_0 \times m}$ is a gadget matrix for which the lattice $\Lambda^\perp(\mathbf{G}_0)$ has a short public basis. A functional secret key for the vector $\vec{y} = (y_1, \dots, y_\ell)^\top$ consists of $\text{dk}_{\vec{y}} = \sum_{i=1}^\ell y_i \cdot \mathbf{s}_i \in \mathbb{Z}^n$ and allows computing $\mathbf{G}_0^\top \cdot (\sum_{i=1}^\ell y_i \cdot \vec{x}_i) + \text{small} \in \mathbb{Z}_q^m$ from $\sum_{i=1}^\ell y_i \cdot \mathbf{C}_i \in \mathbb{Z}_q^m$ and eventually recovering the linear function $\sum_{i=1}^\ell y_i \cdot \vec{x}_i \in \mathbb{Z}^{n_0}$ of $\mathbf{X} = [\vec{x}_1 \mid \dots \mid \vec{x}_\ell] \in \mathbb{Z}_q^{n_0 \times \ell}$.

The construction and proof of [LT19] are not merely obtained by plugging the DPRF of [LST18] into the high-level design principle of [CSG⁺18]. Relying on the DPRF of [LST18] in a modular way seems impossible as it would require a DPRF where partial evaluations are themselves pseudorandom so long as the adversary does not obtain the underlying secret key shares: in the MCFE setting, a challenge ciphertext contains a bunch of partial evaluations (one for each message slot) rather than a threshold recombination of such evaluations. One difficulty is that, in the LWE-based DPRF of [LST18], partial evaluations are not proven pseudorandom: [LST18] only proves – via a deterministic randomness extraction argument – the pseudorandomness of the final PRF value obtained by combining partial evaluations. The proof of [LST18] cannot apply a randomness extractor to individual partial DPRF evaluations as it would destroy their key homomorphic property. Instead of relying on the pseu-

³While their decentralized scheme is only proved secure under static corruptions, its centralized version is proved secure under adaptive corruptions.

dorandomness of partial evaluations, the security proof of [LT19] actually proves a milder indistinguishability property which suffices for their purposes.

The first step is to make sure that all encryption queries will involve a lossy matrix $\mathbf{A}(t)^\top = \mathbf{R}_t \cdot \hat{\mathbf{A}}^\top \cdot \mathbf{C} + \mathbf{E}_t$, for small-norm $\mathbf{R}_t \in \mathbb{Z}^{m \times m}$ and $\mathbf{E}_t \in \mathbb{Z}^{m \times n}$, so that honest senders' ciphertexts are of the form $\mathbf{C}_i = \mathbf{G}_0^\top \cdot \vec{x}_i + \mathbf{R}_t \cdot \hat{\mathbf{A}}^\top \cdot \mathbf{C} \cdot \mathbf{s}_i + \text{noise}$ and thus leak nothing about $\mathbf{s}_i \in \mathbb{Z}^n$ beyond $\mathbf{C} \cdot \mathbf{s}_i \in \mathbb{Z}_q^{n_1}$. The difficulty arises in the challenge queries $(i, t^*, \vec{x}_{0,i}^*, \vec{x}_{1,i}^*)$, where $\mathbf{A}(t^*) \in \mathbb{Z}_q^{n \times m}$ is not a lossy matrix and one must find a way to replace $\mathbf{C}_i^* = \mathbf{G}_0^\top \cdot \vec{x}_{0,i}^* + \mathbf{A}(t^*)^\top \cdot \mathbf{s}_i + \text{noise}$ by $\mathbf{C}_i^* = \mathbf{G}_0^\top \cdot \vec{x}_{1,i}^* + \mathbf{A}(t^*)^\top \cdot \mathbf{s}_i + \text{noise}$ without the adversary noticing. In the DPRF case [LST18], the security proof relies on a deterministic randomness extraction⁴ argument to extract statistically uniform bits from $\lfloor \mathbf{A}(x^*)^\top \cdot \mathbf{s} \rfloor_p$, which has high min-entropy when $\mathbf{A}(x^*)$ is of the form $\mathbf{A} \cdot \mathbf{R}^* + \mathbf{G}$. Here, it is not clear how to apply deterministic extractors in the proof while preserving the functionality of the MCFE scheme.

The solution of [LT19] is to program the public parameters in such a way that, with noticeable probability, the challenge ciphertexts are generated for a matrix $\mathbf{A}(t^*) \in \mathbb{Z}_q^{n \times m}$ of the form

$$\mathbf{A}(t^*)^\top = \mathbf{R}^* \cdot \mathbf{A}^\top + \mathbf{G}_0^\top \cdot \mathbf{V} = \mathbf{R}^* \cdot \hat{\mathbf{A}}^\top \cdot \mathbf{C} + \mathbf{G}_0^\top \cdot \mathbf{V} + \text{noise}, \quad (1)$$

for a statistically random matrix $\mathbf{V} \in \mathbb{Z}_q^{n_0 \times n}$ included in the public parameters. In the proof, the simulator generates a statistically uniform matrix $\mathbf{U} = \begin{bmatrix} \mathbf{V} \\ \mathbf{C} \end{bmatrix}$, where $\mathbf{C} \in \mathbb{Z}_q^{n_1 \times n}$ is used to build the lossy matrix $\mathbf{A}^\top = \hat{\mathbf{A}}^\top \cdot \mathbf{C} + \mathbf{E}$, together with a trapdoor $\mathbf{T}_\mathbf{U}$ for $\Lambda^\perp(\mathbf{U})$. Using $\mathbf{T}_\mathbf{U}$, the simulator can sample a short matrix $\mathbf{T} \in \mathbb{Z}^{n \times n_0}$ satisfying $\mathbf{U} \cdot \mathbf{T} = \begin{bmatrix} \mathbf{I}_{n_0} \\ \mathbf{0} \end{bmatrix} \pmod q$, allowing it to define an alternative secret key $\mathbf{s}'_i = \mathbf{s}_i + \mathbf{T} \cdot (\vec{x}_{0,i}^* - \vec{x}_{1,i}^*) \in \mathbb{Z}^n$. As long as \mathbf{s}_i is sampled from a Gaussian distribution with sufficiently large standard deviation, \mathbf{s}'_i and \mathbf{s}_i are negligibly far apart in terms of statistical distance (as in [Wee14, BBL17], the simulator can guess $\vec{x}_{0,i}^* - \vec{x}_{1,i}^*$ upfront without affecting the polynomial running time of the reduction since the argument is purely statistical). The alternative secret keys $\{\mathbf{s}'_i\}_{i=1}^\ell$ further satisfy $\sum_{i=1}^\ell y_i \cdot \mathbf{s}'_i = \sum_{i=1}^\ell y_i \cdot \mathbf{s}_i$ for all legal functional key queries $\vec{y} = (y_1, \dots, y_\ell)$ made by the adversary. The definition of \mathbf{s}'_i finally ensures that $\mathbf{C} \cdot \mathbf{s}'_i = \mathbf{C} \cdot \mathbf{s}_i \pmod q$, meaning that \mathbf{s}'_i is compatible with all encryption queries for which $\mathbf{A}(t)$ is lossy. From (1), the condition $\mathbf{V} \cdot \mathbf{T} = \mathbf{I}_{n_0} \pmod q$ then implies that the challenge ciphertext can be interpreted as an encryption of $\vec{x}_{1,i}^*$ since $\mathbf{C}_i^* = \mathbf{G}_0^\top \cdot \vec{x}_{1,i}^* + \mathbf{A}(t^*)^\top \cdot \mathbf{s}'_i + \text{noise}$ is statistically close to $\mathbf{C}_i^* = \mathbf{G}_0^\top \cdot \vec{x}_{0,i}^* + \mathbf{A}(t^*)^\top \cdot \mathbf{s}_i + \text{noise}$.

In order to build a DMCFE system, the authors of [LT19] proceed analogously to [CSG⁺18] and combine two instances of their centralized MCFE scheme. The first one is only used to generate partial functional secret keys whereas the second one is used exactly as in the centralized system. As in [CSG⁺18], the DMCFE scheme of [LT19] first has the senders run an interactive protocol allowing them to jointly generate public parameters for the two MCFE instances. At the end of this protocol (which may involve costly multi-party computation operations, but is only executed once), each sender holds an encryption key $\text{ek}_i = (\mathbf{s}_i, \mathbf{t}_i)$ consisting of encryption keys for the two underlying instances. In order to have the i -th sender \mathcal{S}_i generate a partial functional secret key $\text{dk}_{f,i}$ for a vector $\vec{y} = (y_1, \dots, y_\ell)^\top$, the DMCFE scheme of [LT19] exploits the fact that their centralized scheme allows encrypting vectors.

⁴The standard Leftover Hash Lemma cannot be applied since the source $\lfloor \mathbf{A}(x^*)^\top \cdot \mathbf{s} \rfloor_p$ is not guaranteed to be independent of the seed. A deterministic extractor based on k -wise independent functions [Dod00] is thus needed in [LST18].

Namely, the decryptor obtains from \mathcal{S}_i an MCFE encryption of the vector $y_i \cdot \mathbf{s}_i \in \mathbb{Z}^n$ under the encryption key \mathbf{t}_i of the first instance.

Open Problems: Constructing a MCFE scheme that is secure under an LWE assumption with polynomial approximation factor still remains unresolved. Another natural open question is the feasibility of (D)MCFE beyond linear functions under standard assumptions.

4.2 Fully Secure ABE for t -CNF from LWE [Tsa19]

One of the main properties of an ABE scheme is the function class of policies that can be attached to ciphertexts. In fact, ABE was originally suggested as a generalization of identity-based encryption (IBE), in which each ciphertext is destined to a single attribute x (i.e. the policies are point functions). While there are fully secure constructions from bilinear maps for a fairly large class of policies, the situation with lattice-based constructions is less satisfactory and many efforts were made to close this gap. Prior to this work [Tsa19] the only known fully secure lattice construction was for the class of point functions (also known as IBE). In this work, the author constructs for the first time a lattice-based (ciphertext-policy) ABE scheme for the function class t -CNF, which consists of CNF formulas where each clause depends on at most t bits of the input, for any constant t . This class includes NP-verification policies, bit-fixing policies and t -threshold policies.

The starting point of this work is the selectively-secure ABE scheme for circuits of [BGG⁺14]. We describe the difference between full security and selective security. The former is modeled as a game between an adversary \mathcal{A} and a challenger \mathcal{C} as follows. At the beginning of the game, \mathcal{C} publishes the public parameters of the scheme. At any point of the game, \mathcal{A} can query for multiple decryption keys to attributes x of its choice. In the challenge phase, \mathcal{A} chooses a challenge policy f^* and \mathcal{C} returns a ciphertext respective to f^* . The goal of \mathcal{A} is to determine whether this is an encryption of 0 or 1, and the scheme is secure if it cannot do that as long as none of its queried keys x are authorized by f^* . The selective security game is identical, except that \mathcal{A} has to announce the challenge policy f^* before the game begins. In the latter game, the security reduction has the opportunity to generate the public parameters according to f^* . Selective security proofs usually follow a similar structure, where f^* introduces a partitioning of the identity space. The public parameters are generated in the security reduction such that for all x for which $f^*(x) = 0$ (i.e. not authorized by f^*) it is possible to simulate a decryption key, and for all x for which $f^*(x) = 1$, a key for x would allow breaking the hard problem. Since \mathcal{A} can only query for keys of the first type, the reduction can still answer all the queries appropriately.

The main idea that allows to go from selective to adaptive security of ABE in this work relies on the tagging technique of [Gen06], where Gentry presented an adaptively secure IBE scheme as follows. In the real scheme, every ciphertext is associated with a random tag r_{ct} and every key is associated with a random tag r_{sk} . Decryption works as long as the IBE condition is satisfied and $r_{\text{ct}} \neq r_{\text{sk}}$. The probability that decryption fails is negligible since the tags are random. In the security proof, a random degree- Q polynomial P is embedded into the public parameters, such that it is possible to generate a challenge ciphertext respective to any x with the tag $r_{\text{ct}} = P(x)$ and similarly it is possible to generate a key respective to any x with the tag $r_{\text{sk}} = P(x)$. That is, the security reduction can answer any key query and can generate a challenge ciphertext respective to any x . However, if it generates a ciphertext and a key

for the same identity then the decryption fails because they both have the same tag. Recall that in the security game \mathcal{A} is not allowed to query for a challenge and a key respective to the same attribute and therefore it cannot detect that case.

The main idea in this work [Tsa19] is to implement the tagging technique with a *constrained PRF* instead of a random polynomial, which facilitates generalizing beyond point functions. The constrained PRF key for the function f is used as the “tag” r_{ct} of the ciphertext, and decryption with a key for x tagged with r_{sk} is allowed as long as $f(x) = 1$ and $\text{PRF.Eval}_{r_{\text{ct}}} \neq r_{\text{sk}}$. In other words, the ciphertext corresponding to f associates every authorized attribute x with the tag $\text{PRF.Eval}_{r_{\text{ct}}}$. The author uses the lattice techniques of [BGG⁺14] in order to implement this idea. As the constrained PRF needs to satisfy some special structural properties for this construction to work, this work yields a constrained PRF of this form for the function class t -CNF. This specific PRF may be of independent interest, outside the current setup.

- $\text{Setup}(1^\lambda)$: Sample $(\text{cPRF.msk}, \text{cPRF.pp}) \leftarrow \text{cPRF.Setup}(1^\lambda)$ and let $\sigma = \text{cPRF.msk}$. Sample a matrix with its trapdoor $(\mathbf{B}, \mathbf{B}_\tau^{-1}) \leftarrow \text{TrapGen}(1^n, m', q)$. Sample uniformly a matrix $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m \cdot \lambda}$ and a vector $\mathbf{v} \xleftarrow{\$} \mathbb{Z}_q^n$. Output

$$\text{pp} = (\mathbf{B}, \mathbf{A}, \mathbf{v}, \text{cPRF.pp}) \text{ and } \text{msk} = (\mathbf{B}_\tau^{-1}, \sigma).$$

- $\text{Enc}_{\text{pp}}(f, \mu)$: Sample $\text{sk}_f \leftarrow \text{cPRF.KeySim}_{\text{cPRF.pp}}(f)$ and denote $s_f = \text{sk}_f$. Sample $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{e}_0 \xleftarrow{\$} \chi^m$, $\mathbf{e}_1 \xleftarrow{\$} \tilde{\chi}^{m \cdot \ell_f}$, $\mathbf{e}_2 \xleftarrow{\$} \chi$, and output

$$\text{ct} = (s_f, \mathbf{u}_0, \mathbf{u}_1, u_2)$$

such that

$$\mathbf{u}_0 = \mathbf{s}^T \mathbf{B} + \mathbf{e}_0^T, \quad \mathbf{u}_1 = \mathbf{s}^T [\mathbf{A}_f - s_f \otimes \mathbf{G}] + \mathbf{e}_1^T, \quad u_2 = \mathbf{s}^T \mathbf{v} + \mathbf{e}_2 + \mu \lfloor q/2 \rfloor,$$

where $\mathbf{A}_f = \mathbf{A} \mathbf{H}_{\sigma \rightarrow f}$ for $\mathbf{H}_{\sigma \rightarrow f} \leftarrow \text{EvalF}(U_{\sigma \rightarrow f}, \mathbf{A})$.

- $\text{Keygen}_{\text{msk}}(x)$: Compute the matrix $\mathbf{H}_{\sigma \rightarrow x} \leftarrow \text{EvalF}(U_{\sigma \rightarrow x}, \mathbf{A})$ and denote $\mathbf{A}_x = \mathbf{A} \mathbf{H}_{\sigma \rightarrow x}$. Compute $r \leftarrow \text{cPRF.Eval}_\sigma(x)$ and let $I_r : \{0, 1\}^k \rightarrow \{0, 1\}$ be the function that on input r' returns 1 if and only if $r = r'$. Compute $\hat{\mathbf{H}}_r \leftarrow \text{EvalF}(I_r, \mathbf{A}_x)$, denote $\mathbf{A}_{x,r} = \mathbf{A}_x \hat{\mathbf{H}}_r$ and use \mathbf{B}_τ^{-1} to compute $[\mathbf{B} \parallel \mathbf{A}_{x,r}]_\tau^{-1}$. Sample $\mathbf{k} \leftarrow [\mathbf{B} \parallel \mathbf{A}_{x,r}]_\tau^{-1}(\mathbf{v})$ and output $\text{sk}_x = (r, \mathbf{k})$.
- $\text{Dec}_{\text{sk}_x}(\text{ct}, f)$: Parse $\text{sk}_x = (r, \mathbf{k})$ and $\text{ct} = (s_f, \mathbf{u}_0, \mathbf{u}_1, u_2)$. Compute $r' \leftarrow U_{f \rightarrow x}(s_f)$ and if $r = r'$ then abort. Otherwise, compute \mathbf{A}_f and \mathbf{A}_x as in Enc , Keygen respectively, then compute

$$\hat{\mathbf{H}}_{s_f \rightarrow r'} \leftarrow \text{EvalFX}(U_{f \rightarrow x}, s_f, \mathbf{A}_f) \quad \text{and} \quad \hat{\mathbf{H}}_{r,r'} \leftarrow \text{EvalFX}(I_r, r', \mathbf{A}_x).$$

Lastly, compute $u = u_2 - [\mathbf{u}_0 \parallel \mathbf{u}_1 \hat{\mathbf{H}}_{s_f \rightarrow r'} \hat{\mathbf{H}}_{r,r'}] \mathbf{k}$ and output 1 if and only if $|u| \geq q/4$.

Open Problems. Constructing a lattice-based fully-secure ABE for broader classes of functions still remains future work. This can be done either by relaxing the required structural properties of the cPRF (those requirements are inherent in the current [BGG⁺14]-based techniques), or by coming up with other cPRF constructions that satisfy those properties for new function classes.

5 Homomorphic Encryption

Fully-homomorphic encryption (FHE) is an encryption scheme that allows anyone to evaluate any function over encrypted data. Since the breakthrough result of Gentry [Gen09], the development of FHE schemes has seen a rapid surge [vGHV10, BV11, BGV12a, GSW13b, BV14, AP14] and by now FHE has become a well-established cryptographic primitive.

5.1 Rate-1 Fully-Homomorphic Encryption [BDGM19]

We provide an overview of the results and techniques from [BDGM19] that are relevant to PROMETHEUS.⁵ The text in this section may be subject to copyright, please refer to [BDGM19] for information on copyright holders.

Consider a scenario where a party Alice, who has some function f , wants to allow anyone with an input x to learn the evaluation of f on their input, i.e. $f(x)$. The communication complexity of this general problem is well-understood in a model where we do not require any security. In *secure function evaluation* (SFE), we aim to achieve this functionality while providing security to one or both of the parties. Recall that fully-homomorphic encryption immediately gives a two-round protocol for SFE, with communication proportional to the size of the input and of the output, but does not otherwise depend on the size of f . This distinguishing feature is essential for certain applications, such as private information retrieval [CGKS95]. Despite that, known FHE schemes introduce a polynomial blowup factor (in the security parameter) to the length of encrypted messages, thereby affecting the overall communication complexity of such a protocol and making it far from optimal. Thus, a natural approach towards achieving optimal communication would be using FHE scheme with optimal rate, i.e., with a message-to-ciphertext ratio approaching 1, which would immediately give us a general-purpose tool to securely evaluate any function (with sufficiently large inputs and outputs) with asymptotically optimal communication complexity. Given the current state-of-the-art FHE schemes, the only class of functions we can evaluate without communication blowup are linear functions [DJ01].

Motivated by this objective, this work [BDGM19] constructs an optimal-rate FHE scheme. More specifically, the authors show that for any a-priori block size $\ell = \text{poly}$, an FHE scheme can be constructed, where the ciphertext length is at most $\ell + \tau(\lambda)$, where τ is a fixed polynomial that does not depend on ℓ . The scheme is secure under the Learning With Errors (LWE) [Reg05] assumption with polynomial modulus-to-noise ratio.⁶

Apart from its application to SFE with optimal communication, rate-1 FHE essentially improves the communication complexity of any known application of fully-homomorphic encryption. Further, instantiating the generic compiler of Ostrovsky et al. [OPP14] with rate-1 FHE gives the first maliciously circuit-private FHE scheme with optimal rate.

On a technical level, the result relies on the idea of combining a FHE, with a linear decryption algorithm, with a linearly-homomorphic encryption of optimal rate. The hybrid scheme inherits the best of both the worlds and gives us a rate-optimal FHE scheme. The techniques are reminiscent of the chimeric scheme of Gentry and Halevi [GH11], with a new twist on how to encode information without inflating the

⁵Some of the results in the paper use number theoretic assumptions and are not relevant for the project, we will only discuss the results that are based on lattice assumptions.

⁶The modulus-to-noise ratio depends linearly on ℓ .

size of the ciphertexts. Somewhat interestingly, this construction of rate-1 linearly homomorphic encryption from LWE leverages ideas that were originally conceived in the context of spooky FHE [DHRW16], homomorphic secret sharing [BKS19] and private-information retrieval [DGI⁺19].

Technical Outline The focus of this work is on techniques that compress post-evaluation ciphertexts. Compressed ciphertexts can be further expanded via standard bootstrapping techniques.

Schematically, the method for achieving rate-1 FHE is as follows. Consider the “batched-Regev” LWE based encryption scheme (see [PVW08, BGH13] for a detailed description). This scheme has much better rate than “plain” Regev, but the rate is still asymptotically 0 (i.e., $o(1)$). Note that it is possible to convert plain-Regev ciphertexts into batched-Regev, essentially by using the key-switching technique that is frequently used in the FHE literature (see, e.g., [BV11]). The authors further compress batched-Regev ciphertexts in a way that increases the rate to $1 - o(1)$. This is done by combining rounding techniques that appeared previously in the literature [DHRW16, BKS19, DGI⁺19] with new techniques that they develop. These techniques allow them to maintain a high rate, perfect correctness, and modest LWE modulus simultaneously. However, since, in order to apply key-switching, batched-Regev ciphertexts need to be in the non-compressed form, compression techniques are applied only after the switching is complete. This transformation, maintains decryptability but loses homomorphic capabilities, which can be restored using bootstrapping in a generic way.

Leveraging Linear Decryption. The starting point is the observation that in most of the FHE constructions in the literature, decryption (or rather *noisy decryption*) is a linear function $L_c(\mathbf{s})$ in the secret key \mathbf{s} . Typically, for correctly formed ciphertexts c , this linear function satisfies $L_c(\mathbf{s}) = \frac{q}{2} \cdot m + e$, where m is the plaintext and e is a small noise term (say $|e| < B$ for some bound B). One recovers m from $L_c(\mathbf{s})$ via rounding. The choice of the factor $q/2$ is not hardwired into the scheme. It can be provided as an explicit input to the decryption function, in which case this (linear) function looks like

$$L_{\alpha,c}(\mathbf{s}) = \alpha \cdot m + e,$$

for a pre-specified α . The authors call this operation *linear decrypt-and-multiply*.

The authors compose a FHE scheme that has a linear decrypt-and-multiply operation with a rate-1 linearly homomorphic scheme HE to yield a rate-1 FHE. More specifically, they consider a FHE scheme with a linear decrypt-and-multiply operation and where the secret keys are vectors over \mathbb{Z}_q . They further assume that there exists a rate-1 linearly homomorphic scheme HE with plaintext space \mathbb{Z}_q . Then, given a “compression key”, consisting of the encryption $\text{ck} = \text{Enc}(\text{pk}, \mathbf{s})$ of the FHE secret key \mathbf{s} under the linearly homomorphic scheme HE, they compress an FHE ciphertext c encrypting a message $m \in \{0, 1\}$, by transforming c into an encryption of m under HE by homomorphically evaluating the linear function $L_{\alpha,c}(\cdot)$ on ck . In other words, they compute $\text{HE.Eval}(L_{\alpha,c}(\cdot), \text{ck})$. By homomorphic correctness, this results in an encryption of $\alpha \cdot m + e$ under the linearly homomorphic scheme HE.

Rate-1 Linearly Homomorphic Encryption from Standard LWE. One could proceed with the above outline using known rate-1 linearly homomorphic encryption schemes, such as the Damgård-Jurik cryptosystem [DJ01] or a homomorphic variant

of Regev encryption, where the LWE modulus-to-noise ratio is (sub-)exponential [PVW08]. However, in order to get a scheme from the standard LWE assumption, the authors present new constructions of linearly homomorphic encryption schemes from LWE which allow asymptotically optimal ciphertext sizes.

The starting point in this direction is Regev encryption and its variants⁷. Let q be a modulus. A ciphertext c consists of two parts: a vector $c_1 \in \mathbb{Z}_q^n$ and a scalar $c_2 \in \mathbb{Z}_q$. The secret key is a vector $s \in \mathbb{Z}_q^n$. Decryption for this scheme is linear, and it holds that

$$c_2 - s^\top \cdot c_1 = \underbrace{\frac{q}{2} \cdot m + e}_{\hat{m}},$$

where e with $|e| < B$, for some bound B , is a decryption noise term. By computing

$$\lceil \hat{m} \rceil_2 = \lceil \hat{m} \cdot 2/q \rceil = \left\lceil \left(\frac{q}{2} \cdot m + e \right) \cdot 2/q \right\rceil = \lceil m + 2e/q \rceil = m,$$

one can recover the plaintext m , given that $q > 4B$.

The next goal is to shrink the component c_2 of the ciphertext into a single bit. One could think of a solution where the encrypter sends just the rounding of c_2 , i.e. $w = \lceil c_2 \rceil_2$. The decrypter would then recover the message by computing

$$m' = (w - \lceil s^\top c_1 \rceil_2) \bmod 2 = (\lceil c_2 \rceil_2 - \lceil s^\top c_1 \rceil_2) \bmod 2$$

Notice that, since $c_2 - e = s^\top c_1 + m \cdot q/2$, decryption succeeds whenever $\lceil c_2 \rceil_2 = \lceil c_2 - e \rceil_2$. However, although the error term e is small, one could not arbitrarily hope that it does not affect the rounding result.

To guarantee a correct decryption, the decrypter is given an additional value $r \in \mathbb{Z}_q$ such that $c_2 + r \notin [q/4 - B, q/4 + B] \cup [3/4 \cdot q - B, 3/4 \cdot q + B]$. The shrunken ciphertext now consists of $\tilde{c} = (c_1, r, w)$, where $w = \lceil c_2 + r \rceil_2$. Given such a ciphertext \tilde{c} and the secret key s , the decrypter computes

$$m' = (\lceil c_2 + r \rceil_2 - \lceil s^\top c_1 + r \rceil_2) \bmod 2.$$

The careful choice of r insures that $\lceil c_2 + r \rceil_2 = \lceil c_2 + r - e \rceil_2$, and correctness follows.

What is left to be shown is how to amortize the cost of including r by shrinking many c_2 components for the same c_1 . To achieve this, instead of using basic Regev encryption, the authors use *batched* Regev encryption where the ciphertext consists of a vector $c_1 \in \mathbb{Z}_q^n$ and ring elements $c_{2,i} \in \mathbb{Z}_q$, for $i \in [\ell]$, each encrypting a single bit m_i (under a different secret key s_i). They use the same shrinking strategy as above for every $c_{2,i}$. However, now each $c_{2,i}$ imposes a constraint on r for a correct decryption. Fortunately, given that q is sufficiently large one can efficiently compute an r which fulfills all constraints simultaneously. The rate of the resulting scheme is

$$\frac{\ell}{(n+1)\log(q) + \ell} = 1 - \frac{(n+1)\log(q)}{(n+1)\log(q) + \ell}.$$

For $q \approx 4\ell B$ and a sufficiently large $\ell = \Omega(\lambda \cdot (n+1)\log(q)) = \text{poly}$, they achieve rate $1 - O(1/\lambda)$.

Open Problems: This work presents asymptotically optimal schemes, but concrete

⁷While basic Regev encryption is only additively homomorphic, a simple modification transforms it to support evaluation of any linear function.

efficiency remains as a major open problem. In particular it is expected that efficiency gains can be achieved by porting the results to the ring-LWE setting. Steps along these lines, especially in the context of private information retrieval (PIR), were made in an independent work by Gentry and Halevi [GH19]. An additional problem with great importance is to achieve similar rate gains in the context of attribute based encryption (ABE) where the current state of the art is much worse. Note that in the ABE context the rate refers to the ratio between the length of the *attribute* (as opposed to the length of the message) and the length of the ciphertext. An additional line of remaining directions is to explore the applicability of this result in the context of multiparty computation where minimizing the communication complexity also stands as a major challenge.

6 Zero-knowledge

Lattices enable powerful functionalities that are exploited in the extremely active area of lattice-based cryptography. However, they do not easily lend themselves to the realization of certain fundamental tasks, like efficient zero-knowledge proofs. Zero-knowledge protocols ([GMR85]) make it possible to prove properties about certain secret witnesses in order to have users demonstrate their correct behavior while protecting their privacy. For statements proving knowledge of a secret key, efficient solutions are known. See ([MV03],[Lyu08],[KTX08],[LNSW13]). To prove arithmetic relations among committed values, the best known methods rely on the extra algebraic structure offered by ring-LWE or ring-SIS ([LPR10]). However, no truly efficient solution is known from standard (i.e., non-ideal) lattices assumptions. Moreover, the most general zero-knowledge proof techniques can only handle arithmetic circuits in the lattice setting, and adapting them to prove statements over the integers is non-trivial.

6.1 Lattice-Based Zero-Knowledge Arguments [LLNW18]

The problem of constructing zero-knowledge arguments to prove integer-relations among commitments is well-studied. Their importance emanates from the fact that they can be used to prove modular relations when the modulus is unknown at the time of generating the commitment key. The most efficient solutions handling large integers appeal to integer commitments [FO97, DF02] based on hidden-order groups (e.g., RSA groups), which are vulnerable to quantum computing.

In [LLNW18], the authors give statistical zero-knowledge arguments allowing a prover to convince a verifier that x , y and z are commitments to L -bit integers X , Y and Z , respectively that satisfy additive or multiplication relations. Here, the parameter $L = \text{poly}(n)$, where n is the security parameter. Furthermore, the protocol to prove the additive relation $X + Y = Z$ is deployed to prove:

- (i) that the committed integer X belongs to a publicly known range $[\alpha, \beta]$,
- (ii) to prove that X does not belong to a public set, and
- (iii) to prove order relations $Y < X < Z$ between committed integers X , Y , and Z .

This setting differs from the case of arithmetic circuits addressed in [BKLP15] since it deals with proving statements that hold over integers. The non-triviality of this work

stems from the fact that, even in ideal lattices, handling integers of polynomial length L requires working with exponentially large moduli, which affects both the efficiency and the approximation factor of the lattice assumption. In contrast, the protocols described in [LLNW18] use both polynomial moduli and approximation factors.

The protocol emulates integer commitments by means of bit commitments and views integer addition as binary additions with carries. To commit to an L -bit integer X in an all-in-one fashion, generate a KTX [KTX08] commitment

$$\mathbf{c}_x = \sum_{i=0}^{L-1} \mathbf{a}_i \cdot x_i + \mathbf{B} \cdot \mathbf{r} \in \mathbb{Z}_q^n$$

to its binary expansion (x_{L-1}, \dots, x_0) using public matrices $\mathbf{A} = [\mathbf{a}_0 | \dots | \mathbf{a}_{L-1}] \in \mathbb{Z}_q^{n \times L}$ and $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ and randomness $\mathbf{r} \leftarrow U(\{0, 1\}^m)$. The protocol for integer additions has communication cost, roughly, $\tilde{\mathcal{O}}(n + L) \cdot \omega(\log n)$. The prover and the verifier both perform $\mathcal{O}(L)$ simple operations. For proving integer multiplications, this work provides two options. For practically interesting values of L , for instance, $L \leq 8000$, emulate the schoolbook multiplication algorithm by proving L additive relations, and obtain communication cost $\tilde{\mathcal{O}}(n + L^2) \cdot \omega(n)$ as well as computation costs $\mathcal{O}(L^2)$ for both parties. All known methods for proving integer multiplications involve, sometimes implicitly, $\mathcal{O}(L^2)$ computation and/or communication complexities. This work breaks this quadratic barrier and as a theoretical contribution puts forward the first protocol for multiplicative relations that does not incur any quadratic costs. Specifically, by proving in zero-knowledge the correct execution of a Karatsuba multiplication algorithm [KO63], the protocol obtains both computation and communication complexities of order $\mathcal{O}(L^{\log_2 3})$.

Open Problems: One open problem is to achieve similar results using Σ -protocols with large challenge space (possibly using ideal lattices) in order to dispense with the need for $\Omega(n)$ parallel repetitions to achieve negligible soundness error in Stern-like protocols. Indeed, the “Fiat-Shamir with abort” technique does not enable this while guaranteeing the standard knowledge extraction properties.

6.2 Zero-Knowledge Elementary Databases [LNTW19]

Introduced in [MRK03], zero-knowledge sets are protocols that allow a prover P to commit to a finite set S without revealing its size. The commitments are generated efficiently and non-interactively and prove membership or non-membership of certain elements x in the committed set S . *Zero-knowledge elementary databases* (ZK-EDBs) generalize this notion to elementary databases (EDBs). An elementary database D is a partial function; it is a set of key-value pairs (x, y) , where each key x of the universe occurs at most once and thus takes at most one value $y = D(x)$. Keys x that are not in D are assigned $D(x) = \perp$. Every query x obtains a response $D(x)$ and a proof of its correctness. These proofs strictly reveal the value $D(x)$ and nothing else, specially not the size of D . They are sound under the assumption that it is infeasible to find convincing proofs for two distinct values $y \neq y'$ for a given x . Micali et al. [MRK03] described an elegant construction of ZK-EDB based on the discrete logarithm assumption. This was generalized by Chase et al. [CHL⁺05, CHL⁺13] to a general design of ZK-EDBs from a lower-level primitive called mercurial commitment. We describe this briefly later. While efficient and based on standard assumptions, the ZK-EDB realizations from these results have relatively limited expressivity; only simple statements like

“ $x \notin D$ ” or “ $x \in D$ with value $y = D(x)$ ” can be proved. In [LNTW19], the authors further exploit mercurial commitments in order to prove more involved statements like range (of super-polynomial size) queries over keys and values as well as k -nearest neighbour and k -minimum/maximum queries. Their techniques also make it possible to prove membership or non-membership over values, also something that was not known to be possible earlier, without revealing the database size. As a result of independent interest, they construct the ZK-EBD based on a standard quantum-safe assumption in standard (i.e., non-structured) lattices.

There is vast body of work describing and achieving various results for variants of Zero-knowledge databases. In [ORS04], the authors describe protocols for committed databases that can handle orthogonal multi-dimensional range queries, thus allowing for d -dimensional key spaces. While their protocols provide some privacy, they do not hide the database size. There are other efficient constructions known for statistically hiding sets that do not aim at hiding the database size, ([PX09, KZG10]). Whereas work done in [GOSV14] gives black-box constructions of size-hiding database commitments supporting more general queries. Works in ([CFM08, CF13, LY10]) discuss techniques for compressing proofs of membership and non-membership in ZK-EBD’s, under standard number theoretic assumptions. In [GOT15], the authors formalize the notion of ZK-lists and construct size-hiding protocols, in the random oracle, where the prover can demonstrate the order in which elements appear in a committed list.

As mentioned earlier, the partners [LNTW19], construct a non-interactive ZK-EBD protocol in the standard model from mercurial commitments. Briefly, mercurial commitments are commitment schemes that generate commitments in either a hard or soft mode. Commitments in the hard mode satisfy the usual binding property while the ones in soft mode allow the sender to create dummy commitments that do not commit to any particular message. ZK-EDB constructions of ([CHL⁺05, CHL⁺13, MRK03],[LNTW19]) combine mercurial commitments with a Merkle tree, where each internal node contains a mercurial commitment to its two children. The existence of dummy commitments is exactly what allows the sender to commit to the database in polynomial time without revealing its size. The latter is hidden by having a super-polynomial upper bound on the number of leaves in the Merkle tree. Each leaf is assigned to a key x and contains a real commitment to the value $y = D(x)$ and every internal node contains a commitment to its two children. By storing a dummy commitment at the root of each empty subtree, the sender is able to commit to the entire $D = \{(x, y)\}$ in polynomial time. This construction [LNTW19] allows the prover to convincingly answer queries of the form “Give me all database records $(x, y) \in D$ whose keys x lie within the super-polynomial length range $[a_x, b_x]$ ”. Extending this technique also allows the prover to answer queries of the form “Send me all records $(x, y) \in D$ with values y in a super-polynomially large interval $[a_y, b_y]$ ” or prove that “No key x of the database is assigned the value y ” or “ y occurs in D and the corresponding set of keys is $D^{-1}(y)$ ”. All of prover’s responses are polynomially sized. The last two queries are specially interesting as proving them via earlier works would require revealing the size of the database. Integrating the two kinds of queries described above, the construction is also capable of answering range queries over records; given rectangles $[a_x, b_x] \times [a_y, b_y]$ of polynomial width $(b_y - a_y)$ and super-polynomial height $[a_x, b_x]$, the prover yields all the tuples (x, y) in the rectangle. In these instances, the proofs are linear in the size of the width and the number of records in the rectangle. As a special case $[x, x] \times [y, y]$ of range query over records, the construction efficiently proves that specific records (x, y) do not belong to D , which amounts to saying that “if x is in D at all, the corresponding value is not y ”.

Even though these constructions are instantiable with existing mercurial commitments based on standard number theoretic assumptions, the authors in [LNTW19] provide a new construction of trapdoor mercurial commitment (TMC) based on a well-studied assumption in standard (i.e., non-ideal) lattices. In standard lattices, this lattice-based trapdoor mercurial commitment is statistically hiding and computationally binding under the Short-Integer-Solution (SIS) assumption. It performs better than the TMC schemes implied by the generic construction of [CDV06] when the latter is instantiated under the same assumption. It builds on the lattice-based trapdoor commitment (KTX) of Kawachi et al. ([KTX08]) and Micciancio-Peikert trapdoors [MP12].

Open Problems: It would be interesting to see a ZK-EDB protocol that would efficiently answer multi-dimensional range queries without leaking the database size. Even if we disregard quantum security, it would also be interesting to consider the problem of proving statements involving multiple databases committed by the same prover: for example, prove statements about their intersection (e.g., “databases A and B have no more than k elements in common) without even revealing their size.

6.3 RLWE-based Zero-Knowledge Proofs [MM19]

In [Ste93], Stern proposed one of the first post-quantum protocols in his seminal paper for a new identification scheme based on coding theory. His identification protocol was a Zero-Knowledge Proof of Knowledge (ZKPoK) of a solution of an instance of the Syndrome Decoding problem. He gave a 3-move protocol with a soundness error of $2/3$. Many variants and applications have been published since, addressing this lack of efficiency and providing new features. In [KTX08], these techniques are adapted to the setting of lattices still preserving binary secrets. In [JKPT12], Jain *et al.* built a commitment scheme based on the Learning Parity with Noise (LPN) problem, proving knowledge of opening and linear and multiplicative relations between committed messages using 3-move and $2/3$ soundness error Stern-based protocols. Exact Lattice-Based ZKPoK is an active field of research, with very recent efficient constructions for some lattice statements; linear equations with short solutions and matrix-vector relations by Yang et al. [YAZ⁺19], new techniques when a cyclotomic polynomial fully splits in linear factors, by Bootle et al. [BLS19] and new recent Stern-based contributions for proving integer relations [LLNW18] and matrix-vector relations by Libert et al. [LLM⁺19].

In this paper, [MM19], the authors present new and more efficient ways of proving linear and multiplicative relations between elements hidden in lattice-based commitment schemes, without revealing any additional information about the elements themselves. To be precise, this lattice based commitment scheme encodes a message $m \in R_q$ as the coordinates of a point in an ideal lattice defined by $\mathbf{a} \in R_q^k$. To hide this lattice point $\mathbf{a}m$, a RLWE sample $\mathbf{b}r + e$ is added, where $\mathbf{b} \in R_q^k$ generates a lattice distinct from \mathbf{a} , the randomness $r \xleftarrow{\$} R_q$ is chosen uniformly at random and the error term $\mathbf{e} \xleftarrow{\$} \chi^{nk}$ is chosen from an appropriately bounded discrete Gaussian distribution, and $\mathbf{a}m + \mathbf{b}r + \mathbf{e}$ is the commitment outputted. To ascertain the claims about efficiently proving additive and multiplicative relations on messages m , the authors adapt Stern’s protocol to lattices, as in [LNSW13], modify the work of [CVA10] to reduce the soundness error by increasing the number of rounds and improve on

the results of [JKPT12, XXW13] proving linear and multiplicative relations. This 5-move protocol achieves perfect Zero-knowledge with a soundness error $\frac{q+1}{2q}$, which is slightly above $1/2$, as q is usually a very large prime. Further, the commitment scheme used in the protocol is perfectly binding with overwhelming probability over the choice of the public key and is computationally hiding under the RLWE assumption. Multiplication protocol works by masking the original messages and carefully proving that the crossed terms obtained from the multiplication of the masked messages are well-formed.

Compared to previous Stern-based commitment scheme proofs, this construction yields lower computational complexity, improves size of parameters and lowers soundness error for each round. This is justified by the analysis in Fig 1 and Table 3 below. Fig 1 shows the size of the commitment of [XXW13], [BKLP15] and [MM19] for polynomials of degree $n - 1$ where $n = 2^{10}$ and different values of q given by $q \geq n^\gamma$. Then Table 1 shows the communication cost of the proofs for multiplicative relations of the three commitment schemes, where $\kappa = \log(n)$ and $l \cdot d \in \mathcal{O}(n)$, $m/l \in \mathcal{O}(k)$ for a similar level of security.

Figure 1: Commitment's size of Xie *et al.* (—○—), Benhamouda *et al.* (—□—) and Martínez *et al.* (—▲—)




Table 1: Communication cost (in bits)

	Benhamouda <i>et al.</i>	Xie <i>et al.</i>	Martínez <i>et al.</i>
initial com.	-	$md \log^3 q + 2md \log^2 q$	-
round cost	$(8k + 7)n \log q + n/2 + 16\kappa/3 - 8$	$(12\kappa + 2)md \log^3 q + 8ld \log^3 q + \frac{\kappa^2 + 2\kappa + 3}{3}(14md \log q)$	$(3(\kappa + 1)k + 1.5k + 4)n \log q + 6(\kappa + 1)kn + 2 \log q + 1$
aux. com.	1	$3(\log^2 q + 1)$	5
openings	1	$2(\log^2 q + 1)$	3
seeds	-	$2(\kappa \log^2 q + \log^2 q + \kappa)$	$3(\kappa + 1)$

The ideas used in proving multiplicative relations in this work can be easily adapted to any scenario where messages are encoded as RLWE samples. For instance, replacing the lattice-based ZK-PoK construction of [XXW13] in the Attributed Based Sig-

nature scheme for unbounded circuits of [EK18], greatly improves the efficiency of the signature scheme.

Open Problems: This work represents a significant improvement on constructions based on Stern protocol and might be useful in applications that heavily require this kind of proofs, for instance electronic voting. As exemplified above, the ideas in this work are flexible enough to be applied as building blocks for other different constructions besides commitment schemes. Furthermore, it would also be interesting to see an implementation of the protocol presented in this work.

7 Implementation

The development of an efficient quantum order-finding algorithm by Shor [Sho97] invalidated the quantum hardness of factoring and discrete logarithms in Abelian groups, the hardness assumptions in presently used protocols. Since then, there has been a growing effort to develop new encryption algorithms that can resist cryptanalysis using large-scale general quantum computers. Whatever we may think of the timeline or even plausibility of the arrival of general quantum computers, developing quantum-safe cryptography is imperative. From a practical perspective, two crucial requirements are efficiency and ease of deployment of newly proposed schemes. Although, submissions to the NIST process are encouraged to provide optimised software implementations aimed at general purpose microprocessors, implementations of quantum-safe schemes are also required in constrained (often embedded) environments such as micro-controllers or smart cards. These constrained devices are mostly used in embedded applications where low energy consumption, reduced device costs, and other aspects like real-time capabilities are required. For instance in electronic banking, secured identification (passports or national ID cards), authentication, or transport and ticketing applications. The implementation of post-quantum cryptography on constrained devices is an active research area. For instance, [BSJ15], [OPG14], [AJS16], [BBE⁺18], to name a few.

7.1 Implementing RLWE-based Schemes [AHH⁺19]

In this work, [AHH⁺19], the authors re-purpose existing cryptographic coprocessors to feasibly implement lattice-based cryptography in an accelerated manner. Therefore, it would be safe to say that this work can be used by the industry for a possibly smoother migration towards post-quantum cryptography. To be specific, a variant of the Kyber Key Encapsulation Mechanism (KEM) with 161 bits of security is implemented on a commercially available smart card, SLE 78 with 16 Kbyte RAM along with its RSA, AES, and SHA-256 co-processor.

These results are compared with implementations of Kyber and NewHope (Google's first post-quantum at-scale test) on the same target device that does not utilise large integer multiplication, implementations of RSA as well as related work on the co-processor. This comparison makes it evident that lattice-based post-quantum cryptography can be competitive with RSA on contactless high-security 16-bit smart cards with only limited RAM when RSA, AES and SHA-2 co-processors are used. The target chip is equipped with common peripherals (watchdog, timers), internal security functions and encryption procedures, a True Random Number Generator (TRNG), as well as a symmetric co-processor to accelerate AES, a co-processor to compute SHA-256

and an asymmetric co-processor for RSA and ECC acceleration that also allows fast basic long number calculations on integers of size ≈ 2048 bits.

The Kyber KEM scheme is obtained from the RLWE-based KYBER.CPA public encryption scheme via a Fujisaki-Okamoto style transform [HHK17] using a couple of Hash functions. The primary performance stages during implementation are generating the matrix $A \in R_q^{k \times k}$, sampling noise from a distribution on R and evaluating

$$\text{MulAdd}(a(x), b(x), c(x), x^n + 1, q) := a(x) \cdot b(x) + c(x) \pmod{(x^n + 1, q)}$$

in the ring $R := \mathbb{Z}[x] / \langle x^n + 1 \rangle$, for n a power of 2 and a prime modulus q . The efficiency of generating A and sampling noise is directly related to the performance of PRNG() algorithm used to produce large number of pseudorandom data required. It is the speeding up of the polynomial arithmetic that is the central topic of discussion in this work. The original Kyber algorithm uses the Number Theoretic Transform (NTT) for fast polynomial multiplication, which cannot be realized efficiently in the variant considered here. The authors realize this task by using a combination of (a variant of) Kronecker substitution and low-degree polynomial arithmetic. For a large enough ℓ , they define algorithms

$$\begin{aligned} \text{SNORT}(a(x), \ell) &= a(2^\ell) \pmod{2^{2^\ell} + 1} \\ \text{EVAL}(a(2^\ell), b(2^\ell), c(2^\ell), 2^{2^\ell} + 1) &= a(2^\ell) \cdot b(2^\ell) + c(2^\ell) \pmod{2^{2^\ell} + 1} \\ \text{SNEEZE}(D, \ell) &= \{d_i\}_{i=0}^{n-1} \in \mathbb{Z} \\ \text{FINALELL}(\{d_i\}_i) &= d(x) \in R \end{aligned}$$

The table below specifies the number of calls made to these subalgorithm when implementing KYBER.CPA. Here k is the dimension of the secret and the error spaces and MULADDSINGLE denotes big integer multiplication.

	KeyGen	Encrypt	Decrypt
SNORT	$k^2 + k$	$k^2 + 3k + 1$	$2k + 1$
MULADDSINGLE	k^2	$k^2 + k$	k
SNEEZE	k	$k + 1$	1
FINALELL	k	$k + 1$	1

Kronecker substitution is implemented in two ways; using standard Kronecker substitution ($KS1$) together with Karatsuba-based polynomial multiplication and Compact Kronecker ($KS2$) [Har09] using schoolbook-based polynomial multiplication. Implementation with $KS2$ halves the bit size of the output integer at the cost of doubling the number of multiplications. But when compared to RSA, this trade-off seems worthwhile as a decapsulation does 120 modular multiplications of 2049-bit numbers, whereas decrypting 2048-bit RSA requires roughly 3072 multiplications of 1024-bit numbers.

The results show similar performance for the $KS1$ and $KS2$ approach in Kyber.CPA.Imp with a small advantage for $KS2$, whereas Snort for $KS1$ is roughly twice as fast than for $KS2$. However, these conclusions are based on the choice of parameters for this specific co-processor. When compared to a Kyber768 implementation that uses NTT on the SLE78 in software, this approach of using the co-processor to compute the KyberMulAdd gadget provides an advantage, thus showing that NTT may not always be the superior polynomial multiplication algorithm. Finally, this

Kyber variant, that uses the AES co-processor, when run on the target device with an average clock frequency of 50 MHz, uses 72.5 ms to execute Kyber.CPA.Imp.Gen, 94.9 ms for Kyber.CPA.Imp.Enc and 28.4 ms for Kyber.CPA.Imp.Dec. When compared to RSA that does not use CRT, Kyber encryption is slower than RSA encryption but Kyber outperforms RSA in decryption.

Open Problems: These results show that the performance of lattice-based schemes on particular embedded devices highly depends on the speed of the underlying Pseudo Random Number Generator (PRNG). It might be worthwhile to consider constructions that make use of PRNGs based on AES instead of SHA3 due to better availability of (secured) AES hardware acceleration on smart cards or constrained devices in general. The same argument applies to the instantiation of hash functions using SHA-256.

With regard to the optimisation of this particular Kyber implementation, a possible next step is an implementation on an ARM-based smart card or embedded secure element equipped with an ECC/RSA co-processor. On such an architecture the comparison to standard microcontroller-based implementations of PQC would be much easier. Additionally, how much speedup ECC/RSA co-processors will actually provide on ARM platforms equipped with a single-cycle multiplier is still an open question.

In a more general direction it appears interesting to investigate whether a performance advantage can be obtained with schemes specifically designed with the constraints of the big integer multiplier in mind. Although, these schemes use integer sizes too large for direct handling with the current co-processor. In contrast, MLWE-based schemes immediately allow for a piece-wise approach. Another future approach would be a Kyber instantiation with a smaller prime modulus q since in view of the approach described in this work, choice of q is independent of existence of a fast NTT. Moreover, the given results naturally transfer over to the Dilithium signature scheme and an implementation on the SLE 78 is a natural next step. However, parameters have to be adapted for Dilithium. Another interesting question is whether it is possible to efficiently use RSA/ECC co-processors to implement the NTT by treating the big integer multiplier as a vector processor using smart packing of coefficients or a variant of Kronecker substitution.

7.2 A Comparison of the Homomorphic Encryption Libraries [MKLR18]

A homomorphic encryption scheme is one that allows computing on encrypted data without decrypting it first. In fully homomorphic encryption (FHE), it is possible to apply any efficiently computable function to encrypted data. Since the invention of the first fully homomorphic encryption scheme, the landscape of FHE has undergone significant changes. Prototypes demonstrating private health diagnosis, signal processing, genomics statistics, and database queries spur hope on the practical deployment of FHE in the near future.

On the technical side, most of the current applications only consider binary plaintext spaces, and construct binary circuits to compute the desired functions over encrypted data. Although, the existing homomorphic encryption libraries like HElib [HS14, HS15] or SEAL [CLP17] are well adapted to this setting, they also offer the possibility to choose a larger plaintext space, for situations where the function can be evaluated more efficiently when represented by a modular arithmetic circuit. But since that isn't the primary target for these libraries, feasibility remains unclear and unevaluated. Exploring larger plaintext moduli could be interesting for a couple of

more reasons. Firstly, it would make it possible to compute fixed-point high-precision operations over real (truncated or rounded) data. In this case, the modulus bit size is approximately the precision times the multiplicative depth of the circuit. Secondly, outsourced operations over data encrypted using discrete logarithm and factorization-based encryption can be conducted with moduli in the 256-2048 bit range.

In view of this, the authors in [MKLR18] consider the leading homomorphic encryption libraries HELib(-MP), FV-NFLlib and SEAL. They provide a comparative benchmark for large plaintext moduli of up to 2048 bits, analyze their relative performance and compare the impact on the overall performance of the different strategies used in these libraries to handle noise and representation changes. These experiments are conducted on a single core of an Intel(R) Xeon(R) CPU and the parameters are chosen to ensure at least 128 bits of security and at most 12 hour running time. For the readers convenience, HELib implements the BGV homomorphic encryption scheme [BGV12b] and both FV-NFLlib and SEAL implement the Fan-Veracauteren scheme [FV12].

We restrict our exposition to the results obtained when the plaintext modulus p is set at 64 (or 256, or 2048) bit size. The authors provide a detailed analysis of the noise growth with respect to the growing depth for each of the libraries. As the depth grows, it turns out that the growth of the ciphertext modulus size (q), and hence the noise, is comparable in SEAL and FV-NFLlib. Their growth rate is slower than that of HELib by a big multiplicative factor. See figure 2 below.

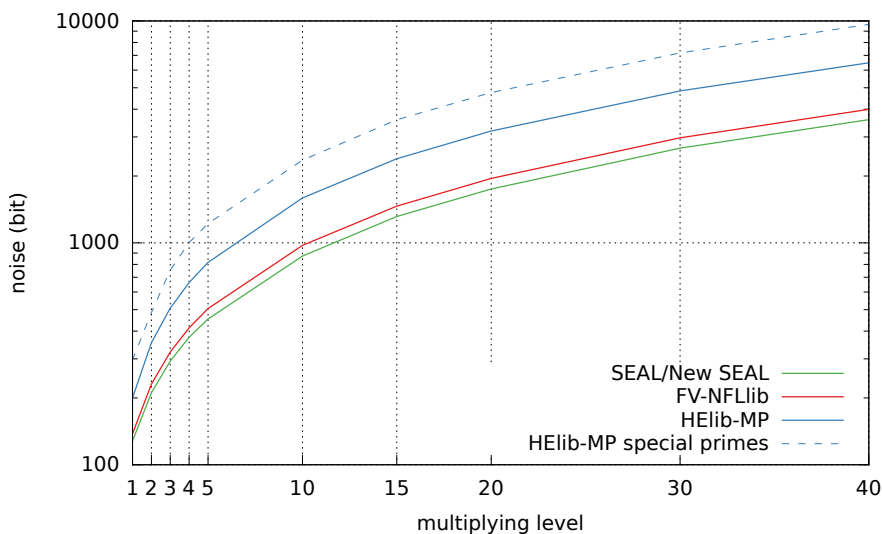


Figure 2: Evolution of the noise as a function of the multiplying depth when $\log p = 64$. For $\log p = 256$ and $\log p = 2048$, results are similar. Results for HELib-MP are given twice, with and without the special primes used in the relinearization operation.

When comparing the computational costs, it is interesting to note that SEAL and FV-NFLlib show almost exactly the same performance results. For HELib, the cost is higher up to a depth of around 40 when $\log p = 64$, and up to a depth of 7, when $\log p = 256$. After that HELib improves significantly and for 2048 bit plaintexts, HELib is better by a factor of 2.5. See figure 3 and 4 below.

Finding the most efficient representation for the plaintext and ciphertext as ring

elements is also a matter of concern. In order to optimize performance, the best would be to switch between DoubleCRT and CRT representations and to avoid lifting coefficients from or projecting them onto their CRT representation, as the latter can become quite expensive.

In conclusion, when $\log p = 60$, FV-NFLlib and SEAL both outperform HELib for

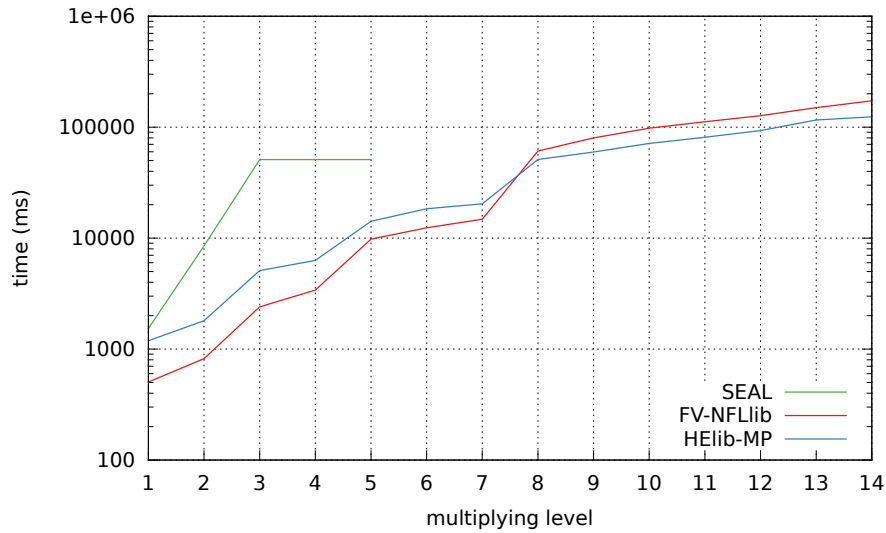


Figure 3: Average time for one multiplication as a function of the multiplying depth when $\log p = 256$.

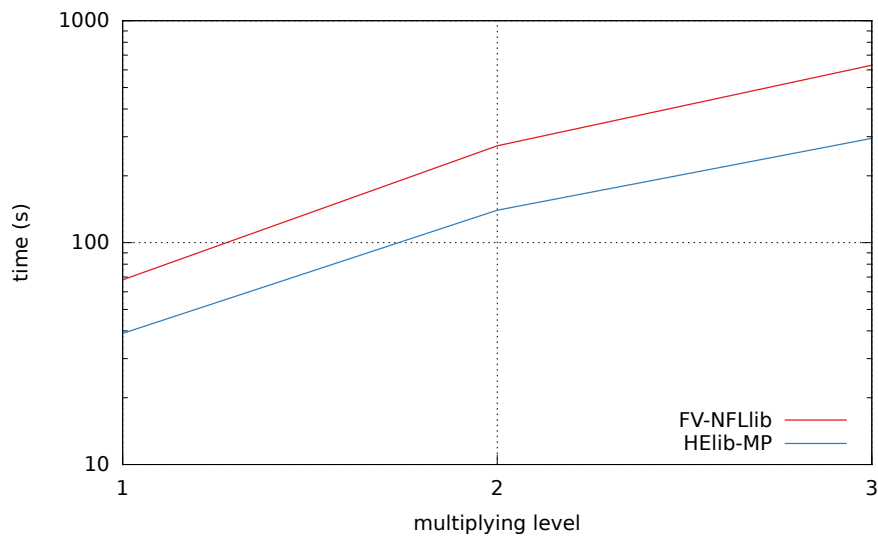


Figure 4: Average time for one multiplication as a function of the multiplying depth when $\log p = 2048$.

all practical values (up to a depth slightly above 40). Given that FV-NFLlib and SEAL result in similar performance and that SEAL is more actively developed and more user friendly, in practice, the natural choice is SEAL. For plaintext modulus above 60, two situations arise. If the plaintext modulus can be factorized into submoduli of 60 bits, then the previous conclusion still applies by using a CRT approach on the plaintext space. If not (for example if it is a cryptographic prime or a hard-to-factor modulus), then SEAL cannot be used and FV-NFLlib gives the best results for $\log q < 2000$, and above that HELib-MP is the best choice.

Open problems: As the BGV homomorphic scheme seems more appropriate for very large moduli, it would be interesting to see a simpler and more optimized implementation based on NFLlib and compare its performance to that of the FV scheme for smaller ciphertext moduli. It would also be interesting to combine the specialized libraries NFLlib and the FullRNS to implement the FV scheme in order to fully exploit its potential.

7.3 On Standardising Sparse-secret LWE Parameter Sets for Homomorphic Encryption [CP19]

The Homomorphic Encryption Security Standard as published by the HomomorphicEncryption.org consortium in 2018, recommends several sets of LWE parameters that can be selected for application in order to achieve a target security level $\lambda \in \{128, 192, 256\}$. All these parameter sets involve a power-of-two dimension $n \leq 2^{15}$, an error distribution of standard deviation $\sigma \approx 3.19$, and a secret whose coefficients are either chosen uniformly from \mathbb{Z}_q , chosen according to the error distribution, or chosen uniformly in $\{-1, 0, 1\}$. However, these sets do not necessarily reflect implementation choices made in the most commonly used homomorphic libraries. For instance, several libraries support dimensions that are not a power of two. Moreover, all known implementations for bootstrapping for the CKKS, BFV and BGV schemes use a sparse secret and a large ring dimension such as $n \in \{2^{16}, 2^{17}\}$. Also, advanced applications such as logistic regression use equally large dimensions. It is for this reason that the recommended parameter sets should be widened to include sets with sparse secrets or larger dimension $n > 2^{15}$.

In this paper, [CP19], the authors explore the security of possible sparse-secret LWE parameter sets, taking into account hybrid attacks, which are often the most competitive in this regime. They present a conservative analysis of the hybrid decoding and the hybrid dual attacks for parameter sets of varying sparsity, with the goal of balancing security requirements with bootstrapping efficiency. They also show how the methodology in the Standard can be easily adapted to support parameter sets with power-of-two dimension $n \geq 2^{16}$.

The security of homomorphic encryption parameter sets is typically determined by considering the best known attacks. Several tools are available to estimate the running time of algorithms for solving LWE. The current version of the Homomorphic Encryption Security Standard uses the LWE Estimator to determine parameters based on the running time of three attacks: *usvp*, *dec* and *dual*. In the setup of LWE with sparse secrets, hybrid attacks ([CHK⁺18], [How07]) are among the most competitive ones. This work analyzes the performance of the following four attacks; *usvp*, *hybrid-dec*, *dual* and *hybrid-dual*. Briefly, *usvp* solves LWE by finding a unique shortest vector in a lattice; *hybrid-dec* or hybrid-decoding solves LWE by decoding a part of the secret and guessing the rest of it appropriately; and *dual* and *hybrid-dual*

solve search-LWE by finding short vectors in a dual lattice, with a guessing component in the latter attack. In the currently standardized LWE parameters for security parameter $\lambda \in \{128, 192, 256\}$, the *usvp* attack performs the best. Exploring parameter sets with sparser secrets, Figure 5 below shows that hybrid attacks are more effective when the secret is chosen from a ternary set of low hamming weight.




Figure 5: An estimate of the ring operations (rop) required to solve the LWE instances parameterised by $n = 1024$, $q = 2^{40}$ and $\sigma \approx 3.2$ and a sparse ternary secret with hamming weight $h \in \{64, 128, 256, 512\}$, using the *usvp*, *dual*, *hybrid-dual* and *hybrid-dec* attacks.

Interestingly, choosing a ternary secret of hamming weight $h = 128$ for the standard set of parameters $(n, \log q, \alpha)$ results in a noticeable security loss. The security drops by approximately 25 bits at the target 128-bit security level, by approximately 50 bits at the target 192-bit security level, and by approximately 85 bits at the target 256-bit security level. One of the main arguments not to standardise sparse secrets is the wider range of attacks that can apply. Moreover, cryptanalysis in this space is very fast-moving; the hybrid-dual attack due to Cheon et al. [CHK⁺18] was only announced recently. This serves to remind us that further away we move from LWE as originally defined, the greater the potential for more efficient attacks. However, the LWE Estimator treats uniform ternary secrets as fixed weight ternary secrets with




Figure 6: Required bitsize $\log q$ to achieve a security level of λ for an LWE instance parameterised by dimension n , modulus q , error standard deviation $\sigma = 3.19$ and a uniform ternary secret with hamming weight $h = \lambda$, under the lattice reduction cost model $T_{\text{BKZ}}(\beta, d) = 2^{0.292\beta + 16.4 + \log(8d)}$. The solid lines represent data points and the dashed lines represent extrapolation to include $n = 65536$ and $n = 131072$.

Hamming weight $h = \lfloor \frac{2n}{3} \rfloor$ and the performance of the hybrid attacks suggests that their future improvements could affect the currently standardized parameters as well. The results further suggest that a Hamming weight $h = \lambda$, for target security level λ , gives a reasonable trade-off between performance and security.

The largest dimension that the existing standard parameter sets allow is $n = 2^{15}$. With current progress in applied homomorphic encryption, the next natural step is therefore to standardize parameter sets for dimension $n = 2^k$, for some $k \geq 16$. In order to further generalize the dimension, i.e. to let n be not a power of two, choosing an error distribution would become more complicated than the convenient coefficient-wise error sampling. As shown in Figure 6, for $n = 2^{16}$, it is straightforward to apply the methodology in the current Homomorphic Encryption Security Standard to use the LWE Estimator to find an appropriate $\log q$ to meet security requirements for fixed $\sigma = 3.19$ and a currently standardised secret distribution.

The same is true in theory for the dimensions $n \geq 2^{17}$, but it becomes cumbersome

some as the LWE Estimator takes hours to run. Therefore, for higher dimensions, results are derived by extrapolating based on known data.

Open Problems: The methodology in the current Standard relies on the LWE Estimator despite a number of limitations. For instance the *dec* estimates output from the Estimator are known to be inaccurate, even though, any standardised parameter sets should be shown to be secure against a state-of-the-art primal decoding attack. Moreover, hybrid attacks are not currently included in the LWE Estimator. Continuing to maintain and improve the LWE Estimator remains important future work.

Another possible direction is revising the methodology used in the Standard: at present, specific sets of parameters are standardised, but it could be more useful to standardise instead the process of using the Estimator itself. This would allow implementors, for example, to use a different error distribution that $\sigma = 3.19$ (with appropriate choices for other parameters) while still conforming to the standard.

8 Conclusion

PROMETHEUS set a goal to produce novel tools for post-quantum cryptographic objects, and WP4 is aimed at developing building blocks that will allow to develop the required cryptographic protocols and applications. In the first half of PROMETHEUS the partners developed a multitude of new tools in the areas that were identified and specified as areas of interest in the grant agreement. These include:

1. Improved quantum resilient versions of basic primitives such as signature schemes, encryption schemes and zero knowledge protocols.
2. Developing novel abilities and functionalities for more advanced cryptographic primitives, that will allow to not only mimic pre-quantum abilities but actually go beyond them. This includes advances in functional encryption, attribute based encryption, homomorphic encryption and advanced signature schemes.
3. To address and promote the practical standpoint of existing and new development, and assess the maturity and readiness for deployment of various solutions.

As outlined throughout this document, the expected progress (as envisioned in the grant agreement) indeed took place via effort by all partners – many times in synergy. This effort produced a high volume of publications in various leading venues, acknowledging the impact that PROMETHEUS is making on research in the studied domains. This state of affairs is in line with our projected goals and appears quite satisfactory (especially given that we are only at the half-way point).

The challenges that WP4 faces in the remainder of the project is to continue developing new cryptographic abilities, and improve the existing ones, to an extent that will allow them to be integrated into the protocols developed in further work packages. This will allow PROMETHEUS and the scientific community as a whole to make use of the advances of PROMETHEUS to benefit society as intended. Specific open problems on each one of the threads of this work package appear in the relevant section.

References

- [ABCP15] Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. Simple functional encryption schemes for inner products. In Jonathan Katz, editor, *Public-Key Cryptography - PKC 2015 - 18th IACR International Conference on Practice and Theory in Public-Key Cryptography, Gaithersburg, MD, USA, March 30 - April 1, 2015, Proceedings*, volume 9020 of *Lecture Notes in Computer Science*, pages 733–751. Springer, 2015.
- [AHH⁺19] Martin R. Albrecht, Christian Hanser, Andrea Höller, Thomas Pöppelmann, Fernando Virdia, and Andreas Wallner. Implementing rlwe-based schemes using an RSA co-processor. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2019(1):169–208, 2019.
- [AJS16] Erdem Alkim, Philipp Jakubeit, and Peter Schwabe. Newhope on ARM cortex-m. In Claude Carlet, M. Anwar Hasan, and Vishal Saraswat, editors, *Security, Privacy, and Applied Cryptography Engineering - 6th International Conference, SPACE 2016, Hyderabad, India, December 14-18, 2016, Proceedings*, volume 10076 of *Lecture Notes in Computer Science*, pages 332–349. Springer, 2016.
- [AKPW13] Joel Alwen, Stephen Krenn, Krzysztof Pietrzak, and Daniel Wichs. Learning with rounding, revisited - new reduction, properties and applications. In *Crypto*, 2013.
- [ALS16] Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In Matthew Robshaw and Jonathan Katz, editors, *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.
- [AP14] Jacob Alperin-Sheriff and Chris Peikert. Faster bootstrapping with polynomial error. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 297–314. Springer, Heidelberg, August 2014.
- [BB04] Dan Boneh and Xavier Boyen. Secure identity based encryption without random oracles. In *Crypto*, 2004.
- [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 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 354–384. Springer, 2018.
- [BBL17] Fabrice Benhamouda, Florian Bourse, and Helger Lipmaa. CCA-secure inner-product functional encryption from projective hash functions. In *PKC*, 2017.

- [BDGM19] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In Dennis Hofheinz and Alon Rosen, editors, *Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part II*, volume 11891 of *Lecture Notes in Computer Science*, pages 407–437. Springer, 2019.
- [BGG⁺14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 533–556. Springer, Heidelberg, May 2014.
- [BGH13] Zvika Brakerski, Craig Gentry, and Shai Halevi. Packed ciphertexts in LWE-based homomorphic encryption. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 1–13. Springer, Heidelberg, February / March 2013.
- [BGV12a] 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.
- [BGV12b] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 309–325. ACM, 2012.
- [BHLY16] Leon Groot Bruinderink, Andreas Hülsing, Tanja Lange, and Yuval Yarom. Flush, gauss, and reload - A cache attack on the BLISS lattice-based signature scheme. In Benedikt Gierlichs and Axel Y. Poschmann, editors, *Cryptographic Hardware and Embedded Systems - CHES 2016 - 18th International Conference, Santa Barbara, CA, USA, August 17-19, 2016, Proceedings*, volume 9813 of *Lecture Notes in Computer Science*, pages 323–345. Springer, 2016.
- [BKLP15] Fabrice Benhamouda, Stephan Krenn, Vadim Lyubashevsky, and Krzysztof Pietrzak. Efficient zero-knowledge proofs for commitments from learning with errors over rings. In Günther Pernul, Peter Y. A. Ryan, and Edgar R. Weippl, editors, *Computer Security - ESORICS 2015 - 20th European Symposium on Research in Computer Security, Vienna, Austria, September 21-25, 2015, Proceedings, Part I*, volume 9326 of *Lecture Notes in Computer Science*, pages 305–325. Springer, 2015.
- [BKPW12] Mihir Bellare, Eike Kiltz, Chirs Peikert, and Brent Waters. Identity-based (lossy) trapdoor functions and applications. In *Eurocrypt*, 2012.
- [BKS19] Elette Boyle, Lisa Kohl, and Peter Scholl. Homomorphic secret sharing from lattices without FHE. In *EUROCRYPT (2)*, volume 11477 of *Lecture Notes in Computer Science*, pages 3–33. Springer, 2019.
- [BLMR13] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In Canetti and Garay [CG13], pages 410–428.

- [BLS19] Jonathan Bootle, Vadim Lyubashevsky, and Gregor Seiler. Algebraic techniques for short(er) exact lattice-based zero-knowledge proofs. In Boldyreva and Micciancio [BM19], pages 176–202.
- [BM19] Alexandra Boldyreva and Daniele Micciancio, editors. *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*, volume 11692 of *Lecture Notes in Computer Science*. Springer, 2019.
- [Bon03] Dan Boneh, editor. *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*. Springer, 2003.
- [BP14] Abhishek Banerjee and Chris Peikert. New and improved key-homomorphic pseudorandom functions. In Garay and Gennaro [GG14], pages 353–370.
- [BPR12] Abhishek Banerjee, Chirs Peikert, and Alon Rosen. Pseudorandom functions and lattices. In *Eurocrypt*, 2012.
- [BSJ15] Ahmad Boorghany, Siavash Bayat Sarmadi, and Rasool Jalili. On constrained implementation of lattice-based cryptographic primitives and schemes on smart cards. *ACM Trans. Embedded Comput. Syst.*, 14(3):42:1–42:25, 2015.
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011.
- [BV14] Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based FHE as secure as PKE. In Moni Naor, editor, *ITCS 2014*, pages 1–12. ACM, January 2014.
- [CDG⁺18] Jérémy Chotard, Edouard Dufour Sans, Romain Gay, Duong Hieu Phan, and David Pointcheval. Decentralized multi-client functional encryption for inner product. In Thomas Peyrin and Steven D. Galbraith, editors, *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part II*, volume 11273 of *Lecture Notes in Computer Science*, pages 703–732. Springer, 2018.
- [CDV06] Dario Catalano, Yevgeniy Dodis, and Ivan Visconti. Mercurial commitments: Minimal assumptions and efficient constructions. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 120–144. Springer, 2006.
- [CF13] Dario Catalano and Dario Fiore. Vector commitments and their applications. In Kurosawa and Hanaoka [KH13], pages 55–72.

- [CFM08] Dario Catalano, Dario Fiore, and Mariagrazia Messina. Zero-knowledge sets with short proofs. In Nigel P. Smart, editor, *Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings*, volume 4965 of *Lecture Notes in Computer Science*, pages 433–450. Springer, 2008.
- [CG13] Ran Canetti and Juan A. Garay, editors. *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*. Springer, 2013.
- [CGKS95] Benny Chor, Oded Goldreich, Eyal Kushilevitz, and Madhu Sudan. Private information retrieval. In *36th FOCS*, pages 41–50. IEEE Computer Society Press, October 1995.
- [CHK⁺18] Jung Hee Cheon, Kyoohyung Han, Andrey Kim, Miran Kim, and Yongsoo Song. Bootstrapping for approximate homomorphic encryption. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 360–384. Springer, 2018.
- [CHL⁺05] Melissa Chase, Alexander Healy, Anna Lysyanskaya, Tal Malkin, and Leonid Reyzin. Mercurial commitments with applications to zero-knowledge sets. In Ronald Cramer, editor, *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings*, volume 3494 of *Lecture Notes in Computer Science*, pages 422–439. Springer, 2005.
- [CHL⁺13] Melissa Chase, Alexander Healy, Anna Lysyanskaya, Tal Malkin, and Leonid Reyzin. Mercurial commitments with applications to zero-knowledge sets. *J. Cryptology*, 26(2):251–279, 2013.
- [CLP17] Hao Chen, Kim Laine, and Rachel Player. Simple encrypted arithmetic library - seal v2.1. In *Financial Cryptography and Data Security*, pages 3–18, 2017.
- [CP19] Benjamin R. Curtis and Rachel Player. On the feasibility and impact of standardising sparse-secret LWE parameter sets for homomorphic encryption. *IACR Cryptology ePrint Archive*, 2019:1148, 2019.
- [CRR02] Suresh Chari, Josyula R. Rao, and Pankaj Rohatgi. Template attacks. In Burton S. Kaliski Jr., Çetin Kaya Koç, and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2002, 4th International Workshop, Redwood Shores, CA, USA, August 13-15, 2002, Revised Papers*, volume 2523 of *Lecture Notes in Computer Science*, pages 13–28. Springer, 2002.
- [CSG⁺18] Jeremy Chotard, Edouard Dufour Sans, Romain Gay, Duong-Hieu Phan, and David Pointcheva. Multi-client functional encryption with repetition for inner product. *Cryptology ePrint Archive: Report 2018/1021*, 2018.

- [CVA10] Pierre-Louis Cayrel, Pascal Véron, and Sidi Mohamed El Yousfi Alaoui. A zero-knowledge identification scheme based on the q -ary syndrome decoding problem. In Alex Biryukov, Guang Gong, and Douglas R. Stinson, editors, *Selected Areas in Cryptography - 17th International Workshop, SAC 2010, Waterloo, Ontario, Canada, August 12-13, 2010, Revised Selected Papers*, volume 6544 of *Lecture Notes in Computer Science*, pages 171–186. Springer, 2010.
- [DDF14] Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: From probing attacks to noisy leakage. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 423–440. Springer, 2014.
- [DDLL13] Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In Canetti and Garay [CG13], pages 40–56.
- [DF02] Ivan Damgård and Eiichiro Fujisaki. A statistically-hiding integer commitment scheme based on groups with hidden order. In Yuliang Zheng, editor, *Advances in Cryptology - ASIACRYPT 2002, 8th International Conference on the Theory and Application of Cryptology and Information Security, Queenstown, New Zealand, December 1-5, 2002, Proceedings*, volume 2501 of *Lecture Notes in Computer Science*, pages 125–142. Springer, 2002.
- [DGI⁺19] Nico Döttling, Sanjam Garg, Yuval Ishai, Giulio Malavolta, Tamer Mour, and Rafail Ostrovsky. Trapdoor hash functions and their applications. CRYPTO, 2019.
- [DHRW16] Yevgeniy Dodis, Shai Halevi, Ron D. Rothblum, and Daniel Wichs. Spooky encryption and its applications. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 93–122. Springer, Heidelberg, August 2016.
- [DJ01] Ivan Damgård and Mats Jurik. A generalisation, a simplification and some applications of Paillier’s probabilistic public-key system. In Kwangjo Kim, editor, *PKC 2001*, volume 1992 of *LNCS*, pages 119–136. Springer, Heidelberg, February 2001.
- [Dod00] Y. Dodis. *Exposure-resilient cryptography*. PhD thesis, MIT, 2000.
- [DT06] Ivan Damgård and Rune Thorbek. Linear integer secret sharing and distributed exponentiation. *IACR Cryptology ePrint Archive*, 2006:44, 2006.
- [EK18] Ali El Kaafarani and Shuichi Katsumata. Attribute-based signatures for unbounded circuits in the ROM and efficient instantiations from lattices. In Michel Abdalla and Ricardo Dahab, editors, *Public-Key Cryptography - PKC 2018 - 21st IACR International Conference on Practice and Theory of Public-Key Cryptography, Rio de Janeiro, Brazil, March 25-29, 2018, Proceedings, Part II*, volume 10770 of *Lecture Notes in Computer Science*, pages 89–119. Springer, 2018.

- [FHPS13] Eduarda Freire, Dennis Hofheinz, Kenneth G. Paterson, and Christoph Striecks. Programmable hash functions in the multilinear setting. In *Crypto*, 2013.
- [FO97] Eiichiro Fujisaki and Tatsuaki Okamoto. Statistical zero knowledge protocols to prove modular polynomial relations. In Burton S. Kaliski Jr., editor, *Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings*, volume 1294 of *Lecture Notes in Computer Science*, pages 16–30. Springer, 1997.
- [FV12] Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. *IACR Cryptology ePrint Archive*, 2012:144, 2012.
- [Gen06] Craig Gentry. Practical identity-based encryption without random oracles. In *Advances in Cryptology – EUROCRYPT '06*, pages 445–464, 2006.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
- [GG14] Juan A. Garay and Rosario Gennaro, editors. *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*. Springer, 2014.
- [GGJS13] Shafi Goldwasser, Vipul Goyal, Abhishek Jain, and Amit Sahai. Multi-input functional encryption. *IACR Cryptology ePrint Archive*, 2013:727, 2013.
- [GH11] Craig Gentry and Shai Halevi. Fully homomorphic encryption without squashing using depth-3 arithmetic circuits. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 107–109. IEEE Computer Society Press, October 2011.
- [GH19] Craig Gentry and Shai Halevi. Compressible fhe with applications to pir. Technical report, 2019. (personal communication).
- [GKL⁺13] S. Dov Gordon, Jonathan Katz, Feng-Hao Liu, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. *IACR Cryptology ePrint Archive*, 2013:774, 2013.
- [GKL⁺14] Samuel Dov Gordon, Jonathan Katz, Feng Hao Liu, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. *Cryptology ePrint Archive: Report 2013/774*, 2014.
- [GKPV10] Shaffi Goldwasser, Yael Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the Learning with Errors assumption. In *ICS*, 2010.
- [GLP12] Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In Emmanuel Prouff and Patrick Schaumont, editors, *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, volume 7428 of *Lecture Notes in Computer Science*, pages 530–547. Springer, 2012.

- [GMR85] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof-systems (extended abstract). In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, pages 291–304. ACM, 1985.
- [GOSV14] Vipul Goyal, Rafail Ostrovsky, Alessandra Scafuro, and Ivan Visconti. Black-box non-black-box zero knowledge. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 515–524. ACM, 2014.
- [GOT15] Esha Ghosh, Olga Ohrimenko, and Roberto Tamassia. Zero-knowledge authenticated order queries and order statistics on a list. In Tal Malkin, Vladimir Kolesnikov, Allison Bishop Lewko, and Michalis Polychronakis, editors, *Applied Cryptography and Network Security - 13th International Conference, ACNS 2015, New York, NY, USA, June 2-5, 2015, Revised Selected Papers*, volume 9092 of *Lecture Notes in Computer Science*, pages 149–171. Springer, 2015.
- [GSW13a] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Crypto*, number 8042 in LNCS, pages 75–92, 2013.
- [GSW13b] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of LNCS, pages 75–92. Springer, Heidelberg, August 2013.
- [Har09] David Harvey. Faster polynomial multiplication via multipoint kronecker substitution. *J. Symb. Comput.*, 44(10):1502–1510, 2009.
- [HHK17] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of LNCS, pages 341–371. Springer, Heidelberg, November 2017.
- [HK08] Dennis Hofheinz and Eike Kiltz. Programmable hash functions and their applications. In *Crypto*, 2008.
- [How07] Nick Howgrave-Graham. A hybrid lattice-reduction and meet-in-the-middle attack against NTRU. In Alfred Menezes, editor, *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*, volume 4622 of *Lecture Notes in Computer Science*, pages 150–169. Springer, 2007.
- [HS14] Shai Halevi and Victor Shoup. Algorithms in helib. In Garay and Gennaro [GG14], pages 554–571.
- [HS15] Shai Halevi and Victor Shoup. Bootstrapping for helib. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 641–670. Springer, 2015.

- [ISW03] Yuval Ishai, Amit Sahai, and David A. Wagner. Private circuits: Securing hardware against probing attacks. In Boneh [Bon03], pages 463–481.
- [JKPT12] Abhishek Jain, Stephan Krenn, Krzysztof Pietrzak, and Aris Tentes. Commitments and efficient zero-knowledge proofs from learning parity with noise. In Xiaoyun Wang and Kazuo Sako, editors, *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 663–680. Springer, 2012.
- [KH13] Kaoru Kurosawa and Goichiro Hanaoka, editors. *Public-Key Cryptography - PKC 2013 - 16th International Conference on Practice and Theory in Public-Key Cryptography, Nara, Japan, February 26 - March 1, 2013. Proceedings*, volume 7778 of *Lecture Notes in Computer Science*. Springer, 2013.
- [KO63] A. Karatsuba and Y. Ofman. Multiplication of many-digital numbers by automatic computers. *Physics-Doklady* 7, 7:595–596, 1963.
- [KTX08] Akinori Kawachi, Keisuke Tanaka, and Keita Xagawa. Concurrently secure identification schemes based on the worst-case hardness of lattice problems. In Josef Pieprzyk, editor, *Advances in Cryptology - ASIACRYPT 2008, 14th International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, Australia, December 7-11, 2008. Proceedings*, volume 5350 of *Lecture Notes in Computer Science*, pages 372–389. Springer, 2008.
- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *Lecture Notes in Computer Science*, pages 177–194. Springer, 2010.
- [LLM⁺19] Benoît Libert, San Ling, Fabrice Mouhartem, Khoa Nguyen, and Huaxiong Wang. Zero-knowledge arguments for matrix-vector relations and lattice-based group encryption. *Theor. Comput. Sci.*, 759:72–97, 2019.
- [LLNW18] Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. Lattice-based zero-knowledge arguments for integer relations. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 700–732. Springer, 2018.
- [LNSW13] San Ling, Khoa Nguyen, Damien Stehlé, and Huaxiong Wang. Improved zero-knowledge proofs of knowledge for the ISIS problem, and applications. In Kurosawa and Hanaoka [KH13], pages 107–124.
- [LNTW19] Benoît Libert, Khoa Nguyen, Benjamin Hong Meng Tan, and Huaxiong Wang. Zero-knowledge elementary databases with more expressive

- queries. In Dongdai Lin and Kazue Sako, editors, *Public-Key Cryptography - PKC 2019 - 22nd IACR International Conference on Practice and Theory of Public-Key Cryptography, Beijing, China, April 14-17, 2019, Proceedings, Part I*, volume 11442 of *Lecture Notes in Computer Science*, pages 255–285. Springer, 2019.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010.
- [LST18] Benoît Libert, Damien Stehlé, and Radu Titiiu. Adaptively secure distributed prfs from LWE. In Amos Beimel and Stefan Dziembowski, editors, *Theory of Cryptography - 16th International Conference, TCC 2018, Panaji, India, November 11-14, 2018, Proceedings, Part II*, volume 11240 of *Lecture Notes in Computer Science*, pages 391–421. Springer, 2018.
- [LT19] Benoît Libert and Radu Titiiu. Multi-client functional encryption for linear functions in the standard model from LWE. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part III*, volume 11923 of *Lecture Notes in Computer Science*, pages 520–551. Springer, 2019.
- [LY10] Benoît Libert and Moti Yung. Concise mercurial vector commitments and independent zero-knowledge sets with short proofs. In Daniele Micciancio, editor, *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, volume 5978 of *Lecture Notes in Computer Science*, pages 499–517. Springer, 2010.
- [Lyu08] Vadim Lyubashevsky. Lattice-based identification schemes secure under active attacks. In Ronald Cramer, editor, *Public Key Cryptography - PKC 2008, 11th International Workshop on Practice and Theory in Public-Key Cryptography, Barcelona, Spain, March 9-12, 2008. Proceedings*, volume 4939 of *Lecture Notes in Computer Science*, pages 162–179. Springer, 2008.
- [Lyu09] Vadim Lyubashevsky. Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *Lecture Notes in Computer Science*, pages 598–616. Springer, 2009.
- [MKLR18] Carlos Aguilar Melchor, Marc-Olivier Killijian, Cédric Lefebvre, and Thomas Ricosset. A comparison of the homomorphic encryption libraries helib, SEAL and fv-nflib. In Jean-Louis Lanet and Cristian Toma, editors, *Innovative Security Solutions for Information Technology and Communications - 11th International Conference, SecITC 2018, Bucharest*,

- Romania, November 8-9, 2018, Revised Selected Papers*, volume 11359 of *Lecture Notes in Computer Science*, pages 425–442. Springer, 2018.
- [MM19] Ramiro Martínez and Paz Morillo. RLWE-based zero-knowledge proofs for linear and multiplicative relations. In *Cryptography and Coding - IMACC, December 16-18, 2019, LNCS 11929 proceedings*, pages 252–277, 2019.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *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 700–718. Springer, 2012.
- [MRK03] Silvio Micali, Michael O. Rabin, and Joe Kilian. Zero-knowledge sets. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 80–91. IEEE Computer Society, 2003.
- [MV03] Daniele Micciancio and Salil P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In Boneh [Bon03], pages 282–298.
- [NPR99] Moni Naor, Benny Pinkas, and Omer Reingold. Distributed pseudo-random functions and KDCs. In *Eurocrypt, 1999*.
- [OPG14] Tobias Oder, Thomas Pöppelmann, and Tim Güneysu. Beyond ECDSA and RSA: lattice-based digital signatures on constrained devices. In *The 51st Annual Design Automation Conference 2014, DAC '14, San Francisco, CA, USA, June 1-5, 2014*, pages 110:1–110:6. ACM, 2014.
- [OPP14] Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 536–553. Springer, Heidelberg, August 2014.
- [ORS04] Rafail Ostrovsky, Charles Rackoff, and Adam D. Smith. Efficient consistency proofs for generalized queries on a committed database. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings*, volume 3142 of *Lecture Notes in Computer Science*, pages 1041–1053. Springer, 2004.
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571. Springer, Heidelberg, August 2008.
- [PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196. ACM, 2008.

- [PX09] Manoj Prabhakaran and Rui Xue. Statistically hiding sets. In Marc Fischlin, editor, *Topics in Cryptology - CT-RSA 2009, The Cryptographers' Track at the RSA Conference 2009, San Francisco, CA, USA, April 20-24, 2009. Proceedings*, volume 5473 of *Lecture Notes in Computer Science*, pages 100–116. Springer, 2009.
- [Reg05] 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.
- [RSV13] Ananth Raghunathan, Gil Segev, and Salil Vadhan. Deterministic public-key encryption for adaptively chosen plaintext distributions. In *Eurocrypt*, 2013.
- [SGGJ⁺14] Samuel Dov Gordon Shaffi Goldwasser, Viapl Goyal, Abhishek Jain, Jonathan Katz, Feng Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. Multi-input functional encryption. In *Eurocrypt*, 2014.
- [Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.
- [Ste93] Jacques Stern. A new identification scheme based on syndrome decoding. In Douglas R. Stinson, editor, *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, volume 773 of *Lecture Notes in Computer Science*, pages 13–21. Springer, 1993.
- [Tsa19] Rotem Tsabary. Fully secure attribute-based encryption for t-cnf from LWE. In Boldyreva and Micciancio [BM19], pages 62–85.
- [VAvHD18] Thijs Veugen, Thomas Attema, Maran van Heesch, and Léo Ducas. Preparing ourselves for the threats of the post-quantum era. *ERCIM News*, 2018(112), 2018.
- [vGHV10] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 24–43. Springer, Heidelberg, May / June 2010.
- [Wee14] Hoeteck Wee. Dual system encryption via predicate encoding. In *TCC*, 2014.
- [XXW13] Xiang Xie, Rui Xue, and Minqian Wang. Zero knowledge proofs from ring-lwe. In Michel Abdalla, Cristina Nita-Rotaru, and Ricardo Dahab, editors, *Cryptology and Network Security - 12th International Conference, CANS 2013, Paraty, Brazil, November 20-22, 2013. Proceedings*, volume 8257 of *Lecture Notes in Computer Science*, pages 57–73. Springer, 2013.
- [YAZ⁺19] Rupeng Yang, Man Ho Au, Zhenfei Zhang, Qiuliang Xu, Zuoxia Yu, and William Whyte. Efficient lattice-based zero-knowledge arguments with standard soundness: Construction and applications. In Boldyreva and Micciancio [BM19], pages 147–175.