

Survey on Fully Homomorphic Encryption, Theory and Applications

Chiara Marcolla, Victor Sucasas, *Member, IEEE*, Marc Manzano, Riccardo Bassoli, *Member, IEEE*,
Frank H.P. Fitzek, *Senior Member, IEEE* and Najwa Aaraj

Abstract—Data privacy concerns are increasing significantly in the context of Internet of Things, cloud services, edge computing, artificial intelligence applications, and other applications enabled by next generation networks. Homomorphic Encryption addresses privacy challenges by enabling multiple operations to be performed on encrypted messages without decryption. This paper comprehensively addresses homomorphic encryption from both theoretical and practical perspectives. The paper delves into the mathematical foundations required to understand fully homomorphic encryption (FHE). It consequently covers design fundamentals and security properties of FHE, and describes the main FHE schemes based on various mathematical problems. On a more practical level, the paper presents a view on privacy-preserving Machine Learning using homomorphic encryption, then surveys FHE at length from an engineering angle, covering the potential application of FHE in fog computing, and cloud computing services. It also provides a comprehensive analysis of existing state-of-the-art FHE libraries and tools, implemented in software and hardware, and the performance thereof.

Index Terms—Fully Homomorphic Encryption, Homomorphic Encryption, Lattices, Neural Networks, Fog Computing, Cloud Computing, IoT.

I. INTRODUCTION

The notion of fully homomorphic encryption, originally called privacy homomorphism, was introduced by Rivest, Adleman and Dertouzos [1] in 1978. For more than 30 years, this concept was considered to be the holy grail of cryptography, until 2009, when Gentry proposed the first fully homomorphic encryption scheme in his PhD thesis [2]. Homomorphic encryption enables operations on plaintexts without decryption. Namely, a set of operations can be performed over ciphertexts such that these operations are reflected as additions and multiplications on the corresponding plaintexts. Thus, homomorphic encryption allows data manipulation in the encrypted domain. This has a tremendous application potential since it allows privacy-preserving data processing, which can be adopted in new emerging fields such as machine

learning, cloud computing, or in the different data processing layers of new generation networks.

Homomorphic encryption schemes that allow one type of operation, or a limited number of operations, have existed for a long time. Some examples are the RSA cryptosystem by Rivest, Shamir and Adleman [3] (1978), encryption scheme of Goldwasser and Micali [4] (1982), ElGamal [5] (1985), Benaloh [6] (1994), Naccache and Stern [7] (1998), Paillier [8] (1999) and Boneh, Goh and Nissim [9] (2005). In particular in [9], the authors proposed the first scheme capable of performing two operations: an arbitrary number of additions and just one multiplication, then again an arbitrary number of additions. However, it was not until 2009, when Gentry [2], [10] proposed the first fully homomorphic encryption (FHE) scheme which supports the evaluation of arbitrary circuits. In his thesis, Gentry not only proposed an FHE scheme, but also provided a method for constructing a general FHE scheme from a scheme with limited but sufficient homomorphic evaluation capacity. Since then, homomorphic encryption has triggered significant interest, and novel constructions on FHE have been proposed following Gentry's idea, being BGV [11], FV [12], TFHE [13], and HEAAN [14] the most representative.

Since the most common scenario for FHE is cloud computing, the majority of research efforts focused on public key FHE schemes. Nevertheless, there are some papers that proposed symmetric key FHE schemes. The symmetric key FHE schemes can be divided into two categories: i) schemes with both symmetric key and public key versions [2], [15]–[17]; ii) and purely symmetric key FHE schemes [18], [19]. Due to their limited applicability for cloud computing, and since the existing ones have no advantage over the public key schemes, FHE schemes that purely enable symmetric key encryption have not gained popularity among the scientific community. This survey only covers public key FHE schemes, but more information on symmetric key schemes can be found in [20]. It is also worth commenting that, in 2011, Rothblum [21] published a method to convert a symmetric key homomorphic encryption scheme with sufficient homomorphic evaluation capacity, into an asymmetric one.

This survey provides a comprehensive vision on homomorphic encryption since its genesis. We extend previous surveys on the topic [20], [22]–[24], to cover the most relevant advances on FHE and its applications. Specifically, the survey is structured as follows: i) Section II provides the preliminaries, containing the required definitions on Number Theory and Probability Theory to understand the constructions of FHE schemes; ii) Section III introduces the mathematical definition

C. Marcolla, V. Sucasas and Najwa Aaraj are with the Technology Innovation Institute, Masdar City Abu Dhabi, United Arab Emirates. (e-mail: {chiara.marcolla,victor.sucasas,najwa.aaraj}@tii.ae).

M. Manzano is with Sandbox Quantum, Palo Alto, CA, US. (e-mail: marc@sandboxquantum.com). M. Manzano is also with the Electronics and Computing Department, Faculty of Engineering, Mondragon Unibertsitatea, Mondragon, Spain.

R. Bassoli and F. H.P. Fitzek are with the Deutsche Telekom Chair of Communication Networks, Institute of Communication Technology, Faculty of Electrical and Computer Engineering, Technische Universität Dresden, Dresden, Germany.

F. H.P. Fitzek is also with Centre for Tactile Internet with Human-in-the-Loop (CeTI) , Cluster of Excellence, Dresden, Germany. (e-mail: {riccardo.bassoli,frank.fitzek}@tu-dresden.de).

of lattices, and the main mathematical problems used in lattice-based cryptography; iii) Section IV defines the concept and the construction of fully homomorphic encryption; iv) Section V caters for an extensive description of FHE schemes in the state of the art, classified according to their generation; and v) Section VI discusses the security of FHE schemes by presenting the different mathematical problems in which they are based. Following a more practical perspective, the survey also describes: i) the application of homomorphic encryption in machine learning in Section VII; ii) the potential adoption of homomorphic encryption for data aggregation in fog computing in Section VIII; iii) previously proposed techniques for the application of homomorphic encryption in cloud computing in Section IX; and iv) homomorphic encryption in practice, describing the current frameworks and libraries, in Section X. Finally, Section XI discusses open challenges in the field and concludes this survey.

II. PRELIMINARIES

This section presents the mathematical foundations and the notation required to understand the developments of homomorphic encryption covered in the next sections.

A. Number Theory

Groups: A group G is a set with an associative operation such that there exists an identity element of G and every element of G has an inverse. If the group operation is commutative, the group is *abelian* (also called *commutative*). For $t \in \mathbb{N}$, a *finite group* \mathbb{Z}_t is the subgroup of positive integers modulo t . This group is also referred to as $\mathbb{Z}/t\mathbb{Z}$ in number theory, and it can be seen as the set of integers in $(-t/2, t/2]$. A multiplicative group \mathbb{Z}_t^\times , for some $t \in \mathbb{N}$, is the set of integers modulo t that are co-prime to t . Note that if t is prime, then \mathbb{Z}_t is the (finite) field \mathbb{F}_t .

Fields: A field is a set with two operations, addition and multiplication. The set is an abelian group under addition with 0 as identity and its nonzero elements are an abelian group under multiplication with 1 as identity. The multiplication is distributive over addition. A *finite field* \mathbb{F}_q is a field with q elements and it exists if and only if q is prime or a prime power. Finally, a *number field* is a vector space with finite dimension over rational numbers \mathbb{Q} and a *cyclotomic field* is a number field obtained by adjoining a complex root of unity to \mathbb{Q} .

Rings: A ring generalizes a field since multiplication does not need to be commutative and some elements do not have a multiplicative inverse. A *quotient ring* $R = \mathbb{Z}[x]/\langle f(x) \rangle$, is a ring of polynomials with integer coefficients modulo the (monic) polynomial $f(x)$. Note that the multiplication between two polynomials in R is a modular multiplication. If the coefficients of the polynomial are in \mathbb{Z}_q (integers modulo q), then we denote it R_q , specifically $R_q = \mathbb{Z}_q[x]/\langle f(x) \rangle$. It is worth pointing out that the ring R is a field if and only if $f(x)$ is an irreducible polynomial over \mathbb{Z} .

Sets: Given a set E , we denote by E^n the set of vectors such that $E^n = \{e = (e_1, \dots, e_n) : e_i \in E\}$. Let E be a commutative ring, then the dot product of two vectors u, v in E^n is defined as $\langle u, v \rangle = \sum_{i=1}^n u_i \cdot v_i$. Similarly, $\mathcal{M}_{h,w}(E)$ denotes the set of matrices of size $h \times w$ with entries in E . Specifically, the term $(\mathbb{Z}_q)^h$ denotes a vector of size h with elements in \mathbb{Z}_q and $(\mathbb{Z}_q)^{h \times w}$ a matrix of size $h \times w$ also with elements in \mathbb{Z}_q .

Ideal: An *ideal* I is a subset of a ring R containing 0 (i.e. the inverse element of the addition) such that the addition of two elements in I is also in I , and the multiplication of an element in I by an element in R is also in I . A *principal* ideal is an ideal generated by one element. In other words, a principal ideal generated by a is the set of multiples of a .

Real Torus: The *real torus* \mathbb{T} is the set \mathbb{R}/\mathbb{Z} of real numbers modulo 1. Note that \mathbb{T} is a group, when using the addition.

Norms: Let x be a vector in E , then we define $\|x\|_\ell := \sqrt[\ell]{\sum_i |x_i|^\ell}$ the ℓ -norm and $\|x\|_\infty := \max_i |x_i|$ the infinity-norm of x , where x_i are the elements in x . We denote by $\|x\|$ the Euclidean norm of the vector x , that is equivalent to the 2-norm. The Euclidean norm can also be referred as the length of a vector. The norms $\|g(x)\|_\ell$ and $\|g(x)\|_\infty$ of a real or integer polynomial $g(x)$ are the norms of its coefficient vector. If $g(x)$ is a polynomial mod $x^n + 1$, we take the norm of its unique representative of degree less or equal than $n - 1$. Moreover, when E is the real torus \mathbb{T} , then the ℓ -norm of $x \in \mathbb{T}^k$ is the ℓ -norm of the representative of x with all coefficients in $(-1/2, 1/2]$. With abuse of notation, we denote it by $\|x\|_\ell$.

B. Probability Theory

Gaussian distribution: The general form of a *normal* (or *Gaussian*) distribution χ for a random variable $x \in \mathbb{R}$ is

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$$

where μ is the center or mean of the distribution and σ is its standard deviation. Note that in the majority of the FHE literature χ is defined as a discrete Gaussian distribution on \mathbb{Z} with centre zero and width parameter αq , denoted by $\mathcal{D}_{\mathbb{Z}, \alpha q}$. Moreover, if σ is bigger than the smoothing parameter $\eta_\epsilon(\mathbb{Z})$ of \mathbb{Z} [25], then the standard deviation of $\mathcal{D}_{\mathbb{Z}, \alpha q}$ is $\sigma = \alpha q / \sqrt{2\pi}$.

Let $R = \mathbb{Z}[x]/\langle f(x) \rangle$ be a polynomial ring. Informally, we denote a *B-bounded* distribution χ over R , if the norm of the coefficients of a polynomial sampled from χ is less than B with overwhelming probability. In general, B is set to be as small as possible while maintaining security (e.g., if χ in R_q , $B \ll q$). In the rest of this work we denote *small element* as any sample from a B -bounded distribution χ .

Negligible and Overwhelming probability: Let $f(\kappa) : \mathbb{N} \rightarrow \mathbb{R}$ be a function where for any possible integer c there exists a value N such that for all $\kappa > N$ it holds that $|f(\kappa)| < \frac{1}{\kappa^c}$. This function is said to be negligible. If the the output of a negligible function is a probability, then the probability is negligible. In cryptography κ is frequently

referred to as security parameter and it represents the length of the secret values. Normally, provably secure schemes are defined by presenting an attack which probability of success is negligible with respect to the security parameter, i.e. it can become arbitrarily small by increasing κ . Analogously, an overwhelming probability is the output of a function $f'(\kappa) = 1 - f(\kappa)$ such that f is a negligible function. Hence, an overwhelming probability can become arbitrarily close to 1 by increasing κ .

III. LATTICES

This section introduces the mathematical definition of lattices, and it also describes the main mathematical problems, on which the security of lattice-based homomorphic encryption schemes relies. The section intends to be self-contained, but it also provides references for interested readers to find full mathematical descriptions and proofs.

A. Definitions

A k -dimensional lattice is a discrete additive subgroup of \mathbb{R}^n . Let $B = (\mathbf{b}_1, \dots, \mathbf{b}_k)$ be linearly independent vectors in \mathbb{R}^n , then we can define the *lattice* $\mathcal{L}(B)$ generated by B as the set of all integer linear combinations of elements of B :

$$\mathcal{L} = \mathcal{L}(B) = \left\{ \sum_{i=1}^k \gamma_i \mathbf{b}_i : \gamma_i \in \mathbb{Z}, \mathbf{b}_i \in B \right\}.$$

The term B is called *base* of the lattice, and the *parallelepiped* associated to the basis B , is defined as

$$\mathcal{P}(B) = \left\{ \sum_{i=1}^k x_i \mathbf{b}_i : x_i \in [-1/2, 1/2) \right\}.$$

The *rank* k of a lattice $\mathcal{L} \subset \mathbb{R}^n$ is the dimension of its linear span, that is, $k = \dim(\text{span}(\mathcal{L}))$. When $k = n$, the lattice is said to be *full rank*. The *volume* of \mathcal{L} (also called *determinant*) is defined as $\text{Vol}(\mathcal{L}) = \sqrt{\det(B^t B)}$. In the special case that \mathcal{L} is a full rank lattice, we have that $\text{Vol}(\mathcal{L}) = |\det(B)|$.

The *dual* of a lattice \mathcal{L} is the set

$$\mathcal{L}^* = \{\mathbf{v} \in \text{span}(\mathcal{L}) : \langle \mathbf{v}, \mathbf{b} \rangle \in \mathbb{Z} \text{ for all } \mathbf{b} \in \mathcal{L}\}.$$

Note that, for any $B \in \mathbb{R}^{n \times n}$, $\mathcal{L}(B)^* = \mathcal{L}((B^{-1})^T)$. From this it follows that $\det(\mathcal{L}^*) = 1/\det(\mathcal{L})$. Moreover, given a matrix $A \in (\mathbb{Z}_q)^{m \times n}$ for some integers q, m, n , we can define two integer q -ary lattices [26],

$$\begin{aligned} \mathcal{L}_q(A) &:= \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x} = A\mathbf{s} \bmod q \text{ for some } \mathbf{s} \in \mathbb{Z}^n\}, \\ \mathcal{L}_q^\perp(A) &:= \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x}A \equiv 0 \bmod q\}. \end{aligned}$$

Note that $\mathcal{L}_q^\perp(A)$ is a *scaled* dual lattice of $\mathcal{L}_q(A)$, namely, $\mathcal{L}_q^\perp(A) = q \cdot \mathcal{L}_q(A)^*$.

It is also worth defining the *ideal lattice* $\mathcal{L}(I)$, which is an integer lattice $\mathcal{L}(B) \subseteq \mathbb{Z}^n$ where $B = \{g \bmod f : g \in I\}$, $I \subseteq \mathbb{Z}[x]/\langle f \rangle$ is an ideal, and f a monic polynomial of degree n .

The *Hermite factor* δ_0^n is defined as

$$\delta_0^n = \|\mathbf{b}_1\|/\text{Vol}(\mathcal{L})^{1/n} \quad (1)$$

where \mathbf{b}_1 is the first basis vector, i.e. the shortest vector, in the base B of the full rank lattice \mathcal{L} . The factor δ_0 is called the *root Hermite factor*. It is worth highlighting that root Hermite factor is an important indicator about the quality of some lattices attacks, as shown in Section VI.

B. Lattice Distance

Most of the known attacks to FHE schemes are based on the notion of *distance*. The concept of distance comes in natural a way. Specifically, for any vector \mathbf{t} in \mathbb{R}^n and any element \mathbf{v} of a lattice \mathcal{L} , the distance between these two vectors is defined as $\text{dist}(\mathbf{t}, \mathbf{v}) = \|\mathbf{t} - \mathbf{v}\|$. Consequently, the minimum distance between \mathbf{t} and any element in \mathcal{L} , is

$$\text{dist}(\mathbf{t}, \mathcal{L}) = \min\{\|\mathbf{t} - \mathbf{v}\| : \mathbf{v} \in \mathcal{L}\}.$$

Let us define the *minimum distance* of lattice \mathcal{L} as $\lambda_1(\mathcal{L})$, which is the length of a shortest non-zero vector in \mathcal{L} , i.e.

$$\lambda_1(\mathcal{L}) = \min\{\|\mathbf{v}\| : \mathbf{v} \in \mathcal{L}, \mathbf{v} \neq 0\}.$$

We can generalize the notion of minimum distance by defining the *i -th successive minimum* $\lambda_i(\mathcal{L})$ as the smallest radius r of a zero-centred ball that contains i (or more) linearly independent lattice points.

A comprehensive explanation of lattices in cryptography can be found in [27], [28] by Peikert and in [26] by Micciancio and Regev.

C. Shortest Vector Problem

The security of many lattice-based FHE schemes relies on the intractability of the shortest vector problem and its variants. The *Shortest Vector Problem* (SVP) consists of finding the shortest non-zero vector in a given lattice. If we restrict the set of input lattices to ideal lattices, then we obtain the *Ideal-SVP* [29]. The relevant variants of SVP problem are defined below:

- *γ -approximate Shortest Vector Problem* (SVP_γ) consists in identifying a vector that is *almost* the shortest vector. Formally, given $\gamma \geq 1$, it consists in finding a non-zero vector $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v}\| \leq \gamma \cdot \lambda_1(\mathcal{L})$.
- *Decisional Shortest Vector Problem* ($\text{GapSVP}_{\gamma,r}$) consists in establishing which given bound for the shortest vector is correct. Specifically, given $\gamma \geq 1$ and $r > 0$, it consists in deciding if either $\lambda_1(\mathcal{L}) \leq r$ or $\lambda_1(\mathcal{L}) \geq \gamma \cdot r$.
- *γ -unique Shortest Vector Problem* (uSVP_γ) consists in finding a shortest non-zero vector in a lattice \mathcal{L} , which satisfies $\lambda_2(\mathcal{L}) > \gamma \lambda_1(\mathcal{L})$ for $\gamma \geq 1$.

Note that if the shortest vector problem is solvable, then the decisional shortest vector problem is also solvable, namely, $\text{GapSVP}_{\gamma,r} \leq \text{SVP}$. Interested readers can find concrete instantiations of the SVP problem in [30].

D. Closest Vector Problem

A generalization of the shortest vector problem is the *Closest Vector Problem* (CVP). In this generalization, a target vector $\mathbf{t} \in \mathbb{R}^n$ is given instead of using the zero vector, and the problem consists of finding a vector in the lattice $\mathbf{v} \in \mathcal{L}$ that is the closest to \mathbf{t} , i.e. $\text{dist}(\mathbf{t}, \mathbf{v}) = \text{dist}(\mathbf{t}, \mathcal{L})$. This problem also has some variants:

- *γ -approximate Closest Vector Problem* (CVP_γ) consists in finding a vector in the lattice that is *almost* a closest to the target vector. Formally, given $\gamma \geq 1$ and $\mathbf{t} \in \mathbb{R}^n$, it consists in finding $\mathbf{v} \in \mathcal{L}$ such that $\text{dist}(\mathbf{t}, \mathbf{v}) \leq \gamma \cdot \text{dist}(\mathbf{t}, \mathcal{L})$.
- *Decisional Closest Vector Problem* ($\text{DCVP}_{\gamma,r}$), given $\gamma \geq 1$, $r > 0$ and $\mathbf{t} \in \mathbb{R}^n$, consists in deciding if either $\text{dist}(\mathbf{t}, \mathcal{L}) \leq r$ or $\text{dist}(\mathbf{t}, \mathcal{L}) \geq \gamma \cdot r$.
- Let $\alpha \leq 1$. *α -Bounded Distance Decoding* (BDD_α) problem consists in, given a target vector \mathbf{t} such that $\text{dist}(\mathbf{t}, \mathcal{L}) < \alpha \lambda_1(\mathcal{L})$, finding a lattice vector $\mathbf{v} \in \mathcal{L}$ closest to \mathbf{t} .

E. Relations Between Shortest and Closest Vector Problems

Kumar and Sivakumar [31] proved that the uSVP_γ problem is NP-hard when $\gamma = 1 + 2^{-n^c}$, for some constant c , whereas Liu, Lyubashevsky and Micciancio [32] showed that the BDD_α problem is NP-hard for $\alpha > 1/\sqrt{2}$. Moreover, Lyubashevsky and Micciancio [33] also proved that $\text{uSVP} \leq \text{BDD}$ and $\text{uSVP} \leq \text{GapSVP}$. More specifically, they proved that for any $\gamma \geq 1$ and $\gamma \leq \text{poly}(n)$, the problem uSVP_γ , $\text{BDD}_{1/\gamma}$ and GapSVP_γ are equivalent up to polynomial approximation factors. Finally, Bai, Stehlé and Wen [34], preprocessing the lattice with Khot's sparsification technique [35], gave a *probabilistic* polynomial-time reduction from $\text{BDD}_{1/(\sqrt{2}\gamma)}$ to uSVP_γ , for any $\gamma > 1$ and γ polynomial in n . The latter restriction was lifted in Wen's PhD thesis [36].

F. Short Integer Solution Problem

The *Short Integer Solution* (SIS) problem was introduced in 1996 by Ajtai [37]. The goal of the SIS problem is to find an integer vector with small norm that is a solution of a given system of integer equations. More specifically, let $q \in \mathbb{Z}$ and let $A \in (\mathbb{Z}_q)^{m \times n}$ be a matrix. Then, given $\beta < q$, the $\text{SIS}_{q,m,\beta}$ problem consists of finding a non-zero vector $\mathbf{x} \in \mathbb{Z}^m$ with $\|\mathbf{x}\| \leq \beta$ such that $\mathbf{x}A \equiv 0 \pmod{q}$. It is worth noting that solving the $\text{SIS}_{q,m,\beta}$ problem is equivalent to finding a vector \mathbf{v} with norm $\|\mathbf{v}\| \leq \beta$ in the scaled dual $\mathcal{L}_q^\perp(A)$ of the lattice $\mathcal{L}_q(A)$. So this problem can be seen as a kind of SVP_γ for this particular family of lattices. Moreover, it is important to highlight that the solution to the SIS problem needs to be bounded on the length, i.e. $\|\mathbf{x}\| \leq \beta$. Without this restriction, it would be easy to find \mathbf{x} with the Gaussian elimination technique.

The ring version of this problem is given by Micciancio in 2002 [38] (extended version in 2007 [39]). Informally, let

$q \in \mathbb{Z}$ and $R_q = \mathbb{Z}_q[x]/\langle f(x) \rangle$ where $f(x) \in \mathbb{Z}[x]$ is a monic polynomial of degree d , and let $\mathbf{a} \in (R_q)^m$ be a vector of m polynomials. Then, given $\beta < q$, the $\text{Ring-SIS}_{q,m,\beta}$ problem [40] consists of finding a non-zero vector of small polynomials $\mathbf{x} \in R^m$ with $\|\mathbf{x}\| \leq \beta$ such that $\mathbf{a}^T \mathbf{x} \equiv 0 \pmod{q}$.

G. Learning With Errors Problem

The key role that lattice-based problems play in cryptography nowadays is especially due to the *Learning With Errors* (LWE) problem and its *Decisional* version, which were introduced by Regev in 2005 [41] (full version [42] in 2009) as an extension of the “learning from parity with error” problem of Blum, Furst, Kearns, and Lipton [43]. Their definitions are provided below:

- Given a vector $\mathbf{b} \in \mathbb{Z}_q^m$ and a matrix $A \in (\mathbb{Z}_q)^{m \times n}$, the LWE problem¹ consists in finding an unknown vector $\mathbf{s} \in \mathbb{Z}_q^n$ such that

$$A\mathbf{s} + \mathbf{e} = \mathbf{b} \pmod{q}$$

where $\mathbf{e} \in \mathbb{Z}_q^m$ is sampled coordinate-wise from an error distribution χ . In other words, the goal is to find a vector $\mathbf{s} \in \mathbb{Z}_q^n$ given a list of $m = n + 1$ *noisy* equations from

$$A_{\mathbf{s},\chi} = \{(\mathbf{a}_i, b_i = \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i) \in \mathbb{Z}_q^n \times \mathbb{Z}_q :$$

$$\mathbf{a}_i \xleftarrow{\$} \mathbb{Z}_q^n, e_i \xleftarrow{\$} \chi\};$$

- the *Decision Learning With Errors* (DLWE) problem consists of distinguishing (with non-negligible advantage) m samples chosen according to $A_{\mathbf{s},\chi}$ (for uniformly random $\mathbf{s} \in \mathbb{Z}_q^n$), from m samples chosen according to the uniform distribution over $\mathbb{Z}_q^n \times \mathbb{Z}_q$.

In [41], Regev reduced the worst case decisional shortest vector GapSVP in a lattice to the LWE problem via a *quantum* reduction. Namely, the paper shows that if it is possible to find an algorithm to solve the LWE problem in polynomial time, then it is also possible to solve the GapSVP problem in polynomial time. Because of that, the security of LWE-based homomorphic encryption schemes is intimately related to lattice problems like SVP and its variants. In the same paper, Regev proved that the LWE problem and its decisional version are computationally equivalent for any prime q that is polynomially-bounded $\text{poly}(n)$. Subsequently, this result was extended to any modulus q in [44]–[48]. In 2009 Peikert [44] (and successively Lyubashevsky and Micciancio [33]) provided a classical reduction from GapSVP problem to LWE with exponential modulus. In [48] Brakerski, Langlois, Peikert, Regev and Stehlé proved that the LWE with *polynomial* modulus is at least as hard as worst-case lattice problems, via a classical reduction.

H. Ring Learning With Errors Problem

The *Ring Learning With Errors* (RLWE) problem, introduced by Stehlé, Steinfeld, Tanaka and Xagawa

¹This instance is also referred to as the *search* version of the LWE problem.

in [29], is the “ring version” of LWE. Specifically, the LWE is in $(\mathbb{Z}_q)^{n+1}$, while in RLWE is in $(R_q)^2$, where $R_q = \mathbb{Z}_q[x]/\langle f(x) \rangle$ where $f(x) \in \mathbb{Z}[x]$ is a monic, irreducible polynomial of degree d and q is a prime.

- The RLWE problem is to discover $s \in R_q$ given access to arbitrarily many independent samples $(a, b = s \cdot a + e) \in R_q \times R_q$, where a is chosen uniformly at random in R_q , and $e \in R_q$ is sampled from an error distribution χ .
- The *Decision Ring Learning With Errors* (DRLWE) problem consists in distinguishing with non-negligible advantage between independent and uniformly random samples in $R_q \times R_q$ and the same number of independent RLWE instances (where $s \in R_q$ is uniformly random).

It is worth highlighting that the secret s can be chosen from the error distribution since, as Applebaum, Cash, Peikert and Sahai [49] proved, it does not affect the hardness of the LWE problem.

Initially, the RLWE problem was called *Ideal*-LWE problem in [29] and, later, its decision version was called *Polynomial Learning With Errors* (PLWE) in [16] by Brakerski and Vaikunthanathan. Also, Lyubashevsky, Peikert and Regev [50] slightly modified the initial definition of RLWE proposed in [29]. Namely, the secret s and the noisy polynomial b are in R_q^\vee , where R^\vee is a *particular* ideal that is dual to R . One of the main contributions in this paper is a search form to the decision reduction.

I. Relations Between LWE, RLWE and Hard Lattice Problems

As explained in Sections III-G and III-H the hardness of LWE and RLWE problems is related to various well-known hard lattice problems [25], [26], [42], [48], [50]. Specifically, the hardness of the RLWE problem is described in [29], where the search variant was proved hard under the ideal-SVP and, in [50] where Lyubashevsky, Peikert and Regev cater for a quantum reduction from the γ -approximate SVP to the RLWE problem and a classical reduction from search to decisional RLWE assumption. Unlike the quantum reduction, which works for any number field and (almost) any modulus, the classical reduction works only for *particular* modulus q and defining polynomial $f(x)$. In [51] Ducas and Durmus partially improved [50] by generalizing over $f(x)$. The restriction on q was lifted by Langlois and Stehlé [52]. Finally, in 2017 Peikert, Regev and Stephens-Davidowitz [53] presented a polynomial-time quantum reduction from worst-case (ideal) lattice problems to decision RLWE which works in any number field and any modulus.

For completeness, we also mention the *General Learning With Errors* (GLWE) problem introduced by Brakerski, Gentry and Vaikunthanathan in [11]. The GLWE problem is the interpolation between the RLWE and LWE problems. Specifically, it consists in finding $s \in R_q^k$ given a list of noisy equations from:

$$\{(a, b = \langle a, s \rangle + e) \in R_q^k \times R_q, \text{ where } a \xleftarrow{\$} R_q^k, e \xleftarrow{\$} \chi\}$$

For $R = \mathbb{Z}$ we have LWE and for $k = 1$ we have RLWE. The GLWE problem was proven to be hard by Langlois and Stehlé in [52], hence generalizing the results of [50]. From a cryptographic point of view, schemes based on the RLWE problem are more computationally efficient and more compact (in terms of ciphertext size) than the ones based on LWE [28]. We refer readers to [28] by Peikert for a comprehensive and more detailed explanation about LWE, RLWE and the hardness of these problems.

IV. FULLY HOMOMORPHIC ENCRYPTION

A *fully homomorphic encryption* scheme can be defined as an encryption scheme where, given some ciphertexts, any operation over the plaintexts can be performed without decryption by manipulating the ciphertexts directly. Such functionality is achievable if and only if addition and multiplication operations can be performed homomorphically, since these two operations constitute a functionally complete set over finite rings. Specifically, any boolean (arithmetic) circuit can be represented using only XOR (addition) and AND (multiplication) gates. In other words, given two ciphertexts $\text{Enc}(x)$ and $\text{Enc}(y)$, where x and y are plaintexts and Enc is the encryption operation, we can obtain the encryption of $x+y$ (or $x \cdot y$) without decrypting $\text{Enc}(x)$ and $\text{Enc}(y)$, by simply adding (or multiplying) these two ciphertexts, and this is sufficient to evaluate any function over encrypted data such that

$$\text{Dec}(\text{Enc}(x) \triangle \text{Enc}(y)) = x \triangle y,$$

where \triangle is the operation sum or product and Dec is the decryption operation.

As shown in Section IV-A, homomorphic encryption is based on probabilistic algorithms. Generally, the encryption procedure adds a random element r which is called *noise* or *error*². The intrinsic feature of FHE is that the error increases any time a homomorphic operation is carried out. So, after a certain number of multiplications (or additions) the ciphertext cannot be decrypted correctly due to the error growth. This problematic limits encryption schemes with homomorphic properties from being fully homomorphic. Only a bounded number of operations can be performed homomorphically before the plaintext cannot be decrypted correctly, hence these schemes are referred to as *somewhat homomorphic*. To overcome this limitation, Gentry [2], [10] introduced a new technique, *bootstrapping*, which is detailed in Section IV-B. This procedure can be used to convert a scheme that is not fully homomorphic into one that is. The method proposed by Gentry can be divided into two steps:

- 1) Creating a *somewhat homomorphic* encryption scheme, that is, a scheme that supports a limited number of homomorphic operations;
- 2) Using *bootstrapping* to reduce the error added during encryption, and making the ciphertext compact.

It is worth commenting that Gentry proposes an additional step, before bootstrapping, to reduce the decryption complexity. This technique, called (*squashing*), and it is used to

²The noise can be added directly as input in the encryption algorithm or included indirectly, e.g. in the public key [54]

express the decryption function as a function with a lower degree. Squashing the decryption circuit is necessary in order to have a bootstrappable decryption circuit in Gentry's scheme. However, unlike bootstrapping, the squashing technique was not generally adopted in subsequent schemes.

We refer the reader to the following interesting overviews on fully homomorphic encryption [22], [23] and to the Homomorphic Encryption Security Standard white paper [55], where the authors provide tables of recommended parameters used for specific FHE schemes at various security levels, considering particular attacks.

A. Definitions and Basic Notions

A public key homomorphic encryption scheme \mathcal{E} [2] is composed of a set of probabilistic polynomial-time (PPT) algorithms (KeyGen, Enc, Dec, Eval) such that:

- The public key-generation algorithm KeyGen takes as input the security parameter λ and outputs the secret key sk , the public key pk and the (public) evaluation key evk , which is needed to perform homomorphic operations over ciphertexts.
- The public encryption algorithm Enc takes as input the public key pk and a message m from the message space. Subsequently, it outputs a ciphertext c .
- The decryption algorithm Dec takes as input the secret key sk and a ciphertext c . Next, it outputs a message m . The algorithm provides as output \perp , if the decryption algorithm cannot successfully recover the encrypted message m .
- The evaluation algorithm Eval takes as input the evaluation key evk , a function f and t -uple of ciphertexts (c_1, \dots, c_t) . It outputs a ciphertext c_f , such that it decrypts to the result of the evaluation of (m_1, \dots, m_t) over f , i.e. $c_f = \text{Eval}_{evk}(f, (c_1, \dots, c_t))$ and $\text{Dec}_{sk}(c_f) = f(m_1, \dots, m_t)$. Note that the ciphertexts c_f and $c \leftarrow \text{Enc}_{pk}(f(m_1, \dots, m_t))$ are equivalent in the sense that they decrypt to the same plaintext, but different in their construction (e.g. they may have different noise levels).

There are two essential characteristics of a homomorphic encryption scheme \mathcal{E} : i) the maximum degree of a function that the scheme supports; ii) the length increase of the ciphertext after each homomorphic operation. The first property defines what functions \mathcal{E} is able to evaluate correctly. Specifically, the scheme \mathcal{E} is \mathcal{F} -homomorphic if it can correctly evaluate any function f in \mathcal{F} , that is, if there exists an evaluation algorithm Eval such that

$$\text{Dec}_{sk}(\text{Eval}_{evk}(f, c_1, \dots, c_t)) = f(m_1, \dots, m_d) \text{ for all } f \in \mathcal{F},$$

where $c_i \leftarrow \text{Enc}_{pk}(m_i)$ for any $i \in \{1, \dots, t\}$. Also, it defines whether an evaluated ciphertext c_f , namely an output of Eval, can be used as an input of the evaluation algorithm. Specifically, in a *multi-hop homomorphic* scheme [56], a sequence

of any (polynomial) i functions can be homomorphically evaluated one by one on a ciphertext c produced by encrypting a message m . The second property refers to the ciphertext expansion, i.e. how much the ciphertext bit-length grows after each evaluation. In the sense, if the bound of the bit-length growth is independent of the complexity of f it is called *compact* [56].

Depending on the previous notions, we have different definitions of homomorphic encryption schemes:

- *Fully Homomorphic Encryption* scheme \mathcal{E} is an encryption scheme where the ciphertexts are compact, and the scheme is \mathcal{F} -homomorphic, where \mathcal{F} is the set of all the (efficiently computable) functions [57].
- *Leveled fully homomorphic* scheme is a scheme that supports the evaluation of specific depth circuits. More formally, it is \mathcal{F} -homomorphic, where \mathcal{F} is the set of all functions of some specific degree and the bound over the length of the ciphertext is independent of such degree (i.e. it is compact) [11]. In this type of schemes the depth is treated as a setup parameter of the scheme and can adopt any value. It is worth commenting that if the degree is bounded to a maximum value L then the scheme is called *L-Leveled fully homomorphic* scheme. Note that L-leveled schemes may not be compact.
- *Somewhat Homomorphic Encryption* (SHE) scheme is a scheme that is \mathcal{F} -homomorphic for a limited class \mathcal{F} , e.g. capable of evaluating "low-degree" multivariate polynomials homomorphically [11]. Similarly to L-leveled schemes, in an SHE scheme the compactness of ciphertexts could be violated.

It is worth highlighting that, as mentioned in Section IV, an SHE scheme (and a leveled fully homomorphic scheme), with sufficient homomorphic evaluation capacity, can be transformed into an FHE scheme by using the bootstrapping technique.

B. Bootstrapping

Bootstrapping is a technique to decrease the error of the ciphertext, proposed by Gentry in [2], [10]. Essentially, it consists of a re-encryption procedure of a ciphertext c to *refresh* it, i.e. encrypt it again under another key obtaining a new ciphertext (for the same plaintext) but with a smaller error.

Let us consider a somewhat homomorphic encryption scheme \mathcal{E} , two pairs of keys (sk_1, pk_1) and (sk_2, pk_2) , the encryption algorithm Enc and the ciphertext $c = \text{Enc}_{pk_1}(m)$ that encrypts m under pk_1 . The procedure to refresh c is conducted in three steps as follows:

- Encrypting the secret key sk_1 under pk_2 : $\text{Enc}_{pk_2}(sk_1) \leftarrow \overline{sk_1}$. The encryption of sk_1 may require its bit decomposition and thus produce many ciphertexts.

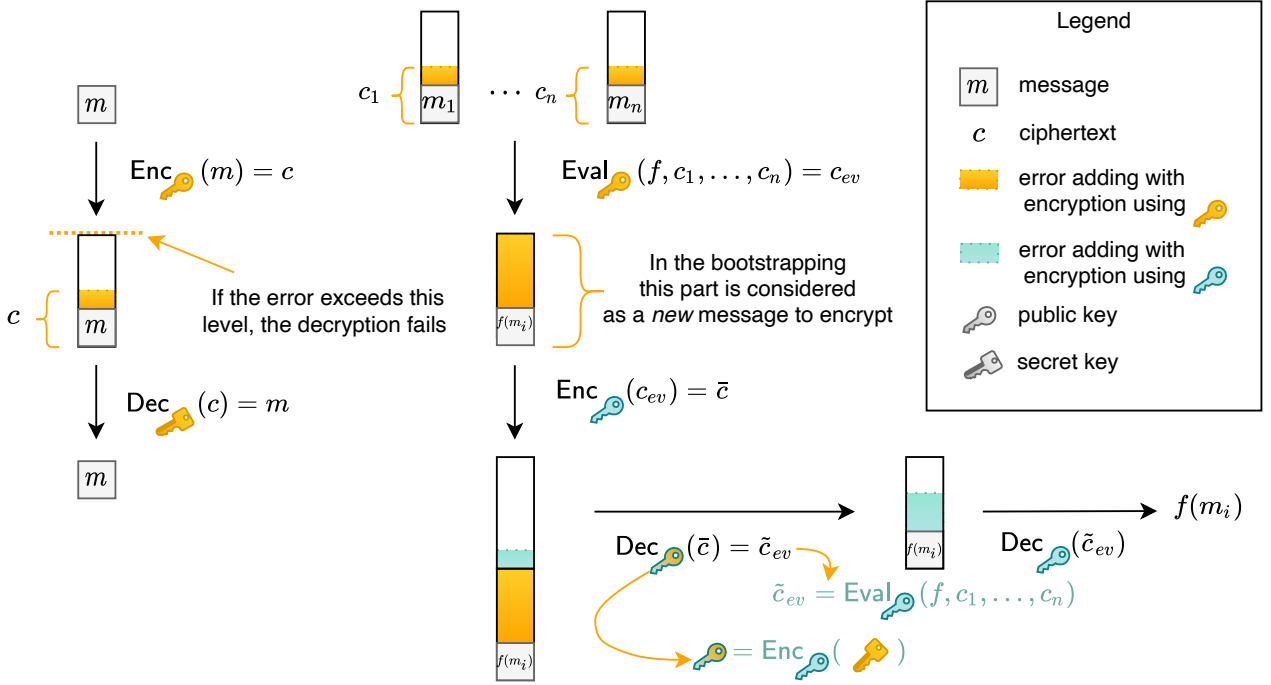


Fig. 1. Bootstrapping technique.

- Encrypting the ciphertext c under pk_2 : $Enc_{pk_2}(c) = \bar{c}$. In most schemes this step is essentially vacuous, in the sense that $Enc_{pk_2}(c)$ is obtained by using null randomness, i.e. viewing a plaintext directly as a ciphertext (with proper padding/scaling)
- Decrypting homomorphically the *new* ciphertext using the encrypted secret key: $Dec_{\bar{sk}_1}(\bar{c})$. In this way, an encryption of the same message under the second public key $Enc_{pk_2}(m)$ is obtained. Namely:

$$\begin{aligned} Eval_{evk}(Dec, \bar{c}, \bar{sk}_1) &= \\ Eval_{evk}(Dec, Enc_{pk_2}(c), Enc_{pk_2}(sk_1)) &= \\ \widehat{Enc}_{pk_2}(Dec_{sk_1}(c)) &= \widehat{Enc}_{pk_2}(m) \end{aligned} \quad (2)$$

the value $\widehat{Enc}_{pk_2}(m)$ is equivalent to $Enc_{pk_2}(m)$ in the sense that both decrypt to m , i.e

$$Dec_{sk_2}(\widehat{Enc}_{pk_2}(m)) = Dec_{sk_2}(Enc_{pk_2}(m)) = m$$

The objective of bootstrapping is to reduce the error of the ciphertext. Conceptually, bootstrapping applies the decryption function and simultaneously performs a second encryption. These operations produce a new ciphertext. This new ciphertext contains the error of a new encryption plus the error increase resulting from the homomorphic evaluation of the decryption circuit. Hence, the error of the obtained ciphertext is higher than a fresh ciphertext (obtained with the encryption algorithm) but lower than a ciphertext obtained after homomorphic evaluating functions with higher depth than the decryption circuit. The idea of bootstrapping is illustrated in Figure 1. Bootstrapping can be applied to any ciphertext, hence it can be used after several homomorphic evaluations to reduce the error

level. This enables the homomorphic evaluation of functions with arbitrarily large depth, but requires the decryption circuit to be bootstrappable.

The cornerstone result obtained by Gentry is the proof that to construct an FHE scheme suffices to construct a scheme that is capable of evaluating only a particular set of small degree functions, i.e. an SHE scheme, which security holds when the encryption of the secret key is published. If that is the case, then bootstrapping can be applied (and squashing if necessary) to obtain an FHE scheme. It is worth highlighting that bootstrapping is the only known way to obtain fully homomorphic encryption schemes. Unfortunately, bootstrapping is computationally complex and requires a large memory space.

C. Security Properties

A homomorphic encryption scheme must be semantically secure, but optionally it can also be function or circuit private. The scheme is secure if and only if it is semantically secure. Semantic security is formally captured by the concept of indistinguishability under chosen-plaintext attack (IND-CPA security), where an attacker can obtain encryptions of arbitrary plaintexts, but it cannot decrypt arbitrary ciphertexts (note that the encryption algorithm and the encryption key pk are public, whereas the decryption key sk is not). If we restrict the message space to $\{0, 1\}$, then:

Let $\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval})$ be a public homomorphic scheme, and m_b a message with $\{0, 1\}$ as the message space. Let us define an adversary \mathcal{A} that knows the evaluation key evk and the public key pk and is given an encryption $Enc_{pk}(m)$ for $m \in \{0, 1\}$. \mathcal{A} can make queries to the encryption oracle. After a polynomial number of queries, \mathcal{A} tries to guess whether




Fig. 2. Timeline of the main FHE schemes.

$m = 0$ or $m = 1$. Then, the scheme is IND-CPA secure if for an efficient adversary \mathcal{A} , it holds that:

$$\left| \Pr[\mathcal{A}(\text{pk}, \text{evk}, \text{Enc}_{\text{pk}}(0)) = 1] - \Pr[\mathcal{A}(\text{pk}, \text{evk}, \text{Enc}_{\text{pk}}(1)) = 1] \right| = \text{negl}(\lambda)$$

where $(\text{sk}, \text{pk}, \text{evk}) \leftarrow \text{KeyGen}(\lambda)$.

The definition above expresses the fact that the adversary is not able to tell apart encryptions of 0 from encryptions of 1 with non-negligible probability. It is worth noting that IND-CPA security is only achievable if the encryption scheme randomizes the ciphertexts. If there is no randomization, the adversary can encrypt messages, and then compare them with the received ciphertext $\text{Enc}_{\text{pk}}(m)$.

Finally, a homomorphic encryption scheme that remains IND-CPA secure against adversaries who observe an encryption of the scheme's secret key under its public key is called *circular* secure. Current constructions of fully homomorphic encryption schemes require an encryption of the secret key to be bootstrappable, hence all known FHE constructions are circular secure. Optionally, the homomorphic encryption scheme can be *function private*. That is, a ciphertext that has been homomorphically evaluated over a function f and does not reveal any information about f , beyond the outputs for the queried inputs. Function privacy is the relaxed version of the original *circuit privacy* [2], which requires that the evaluated ciphertext is statistically indistinguishable from a fresh ciphertext [58]. Note that for a scheme to be circuit or function private, the property has to hold even against an adversary that knows the secret key and can decrypt any ciphertext.

V. FULLY HOMOMORPHIC ENCRYPTION SCHEMES

Gentry's seminal works [2], [10] paved the way for novel FHE schemes in four main research branches: i) the schemes based on ideal lattices (Section V-A); ii) the schemes over integers based on the *Approximate - Greatest Common Divisor* (AGCD) problem (Section V-B); iii) the schemes based on the LWE problem and on its ring version (Sections V-C, V-E

and V-F); and iv) the schemes based on NTRU (Section V-D). Additionally, other research works have proposed schemes based on other mathematical problems (Section V-H). Figure 2 describes the timeline of the main FHE schemes.

A. First Generation: FHE based on ideal lattices

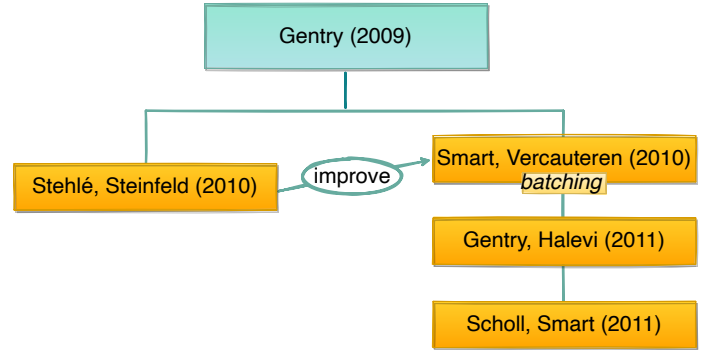


Fig. 3. Main FHE schemes based on ideal lattices.

The first fully homomorphic encryption scheme presented by Gentry [2], [10] is based on ideal lattices. Figure 3 describes the main schemes based on ideal lattices.

Following the notation of [59], we describe one of the constructions presented by Gentry. This scheme uses the integer sublattice $\mathcal{L}(I) \subseteq \mathbb{Z}^m$ defined by the ideal I . The encryption and decryption algorithms work as follows:

- *encryption*: the message $m \in \mathbb{F}_2$ (the message is a single bit, $m \in \{0, 1\}$) is encoded into a point $\mathbf{a} = m + 2\mathbf{e}$ in \mathbb{R}^d , where \mathbf{e} is a small error, namely a random vector with coefficients in $\{0, \pm 1\}$, where the values ± 1 are taken with equal probability. The ciphertext \mathbf{c} is the translation of \mathbf{a} into the parallelepiped $\mathcal{P}(B_{\text{pk}})$ where B_{pk} is the public key. Namely, $\mathbf{c} = \mathbf{a} - (\lceil \mathbf{a} B_{\text{pk}}^{-1} \rceil B_{\text{pk}})$, where $\lceil \cdot \rceil$ denotes rounding to the nearest integer.
- *decryption*: computes $\mathbf{a}' = \mathbf{c} - (\lceil \mathbf{c} B_{\text{sk}}^{-1} \rceil B_{\text{sk}})$, which is the translation of \mathbf{c} into the parallelepiped $\mathcal{P}(B_{\text{sk}})$ where B_{sk} is the secret key, and then outputs $\mathbf{a}' \bmod 2$ as the

decrypted plaintext.

Both the public and secret keys, B_{pk} and B_{sk} , are bases of the ideal lattice $\mathcal{L}(I)$. However, the public key is a bad basis for I , in the sense that it is formed by skewed vectors, whereas the secret key is a good basis because it is formed by orthogonal vectors (see Figure 4).




Fig. 4. Good and bad bases for an ideal lattice.

Unfortunately, the scheme that we just described is not *bootstrappable* due to the complexity of the decryption algorithm (i.e. it cannot be evaluated homomorphically). To solve this problem, Gentry proposed a method to *squash* the decryption function of an SHE scheme. This method consists of transforming the original SHE scheme \mathcal{E} into another scheme \mathcal{E}^* with the same homomorphic capacity but with a simpler decryption function that allows the bootstrapping. To reduce the decryption algorithm's complexity, Gentry proved [2], [10] that it is enough to add to the evaluation key some “extra information” about the secret key. Such extra information consists of a set of vectors $\mathcal{S} = \{s_i \mid i = 1, \dots, S\}$ from which a subset T is derived. The secret key sk is the sum of elements of T , and the public information included in the evaluation key is the set \mathcal{S} . The security of this new scheme \mathcal{E}^* is based on fact that set T is sparse and secret such that the *Sparse Subset Sum Problem* (SSSP) applies. SSSP consists in verifying, given a set of n integers $S = \{a_1, \dots, a_n\} \subseteq \mathbb{Z}$, whether a sparse subset of S exists such that $\sum_{i \in I} a_i = 0$, where $I \subseteq \{1, \dots, n\}$. The scheme is circularly secure, and its security is based on three mathematical problems: i) the *Sparse Subset Sum Problem*; ii) the *Bounded Distance Decoding Problem* (Section III-D); iii) the *Ideal-Shortest Vector Problem* (Section III-C).

Gentry's scheme was initially implemented by Smart and Vercauteren [60] who used principal ideal lattices and introduced the *batching technique*. The batched version of a scheme enables the packing of a vector of ℓ plaintexts to be encrypted in a single ciphertext using the Chinese Remainder Theorem (CRT). This technique permits processing several messages

simultaneously. The Smart-Vercauteren implementation was improved in 2011 by Gentry and Halevi [61], who also simplified the *squashing procedure*. Successively, Scholl and Smart [62] ameliorated the Gentry-Halevi technique providing a generalization over any cyclotomic field. Stehlé and Steinfeld in [63] reduced the bit complexity (i.e. quantity of operations per bit) for refreshing the ciphertext. Their technique can be applied to different FHE schemes, such as Gentry [2] and Smart and Vercauteren [60].

A drawback of ideal lattice-based FHE schemes is that they are based on mathematical constructions that are difficult to implement efficiently. Also, a vulnerability in schemes using principal ideals was found by Cramer, Ducas, Peikert and Regev [64] in 2016. Specifically, a key-recovery attack for cryptographic constructions based on principal ideal lattices is possible, given a quantum polynomial-time or classical $2^{n^{2/3-\epsilon}}$ -time algorithm for finding the short generator of the principal ideal problem.

B. First Generation: FHE based on the AGCD problem

A new (and simpler than ideal lattice-based) family of FHE schemes dawned in 2010 thanks to van Dijk, Gentry, Halevi and Vaikuntanathan, who introduced a fully homomorphic encryption scheme over integers [17]. The basic construction of the DGHV scheme is the following:

- *key generation*: outputs the secret key p , i.e. an odd random integer, and the public key (x_0, \dots, x_n) where x_0 is odd and $x_0 > x_i \forall i$, where $x_i = pq_i + r_i$ with q_i, r_i random integers.
- *encryption*: the message $m \in \mathbb{F}_2$ is encoded into the ciphertext $c = (m + 2r + 2\sum_{i \in S} x_i) \bmod x_0$, where r is a random integer and S is a random subset of $\{1, \dots, n\}$.
- *decryption*: computes $(c \bmod p) \bmod 2$.

The security of this scheme is based on SSSP, circular securitythe Approximate - Greatest Common Divisor (AGCD) problem which consists of finding the “common near divisor” p , given a set of integers $\{x_0, \dots, x_n\} \in \mathbb{Z}$, all randomly chosen and close to multiples of a large integer p .

The main drawbacks of the DGHV scheme are its high computational complexity and large public key size. Several optimizations and implementations have been proposed, as depicted in Figure 5. Specifically, in [65], Coron, Mandal, Naccache and Tibouchi reduced the elements of the public key that have to be stored ($2\sqrt{n}$ instead of n elements), since the rest of the key elements can be recovered. This optimization only requires a slight modification of the encryption procedure. In [66] Chan and Nguyen presented new algorithms to solve the AGCD problem, which are exponentially faster than the previous one. As a consequence, they proved that the scheme in [65] achieves lower security level than initially claimed. Later Coron, Naccache and Tibouchi [67] further reduced the size of the public key using the *modulo switching* technique. The security of both these schemes is based on the AGCD

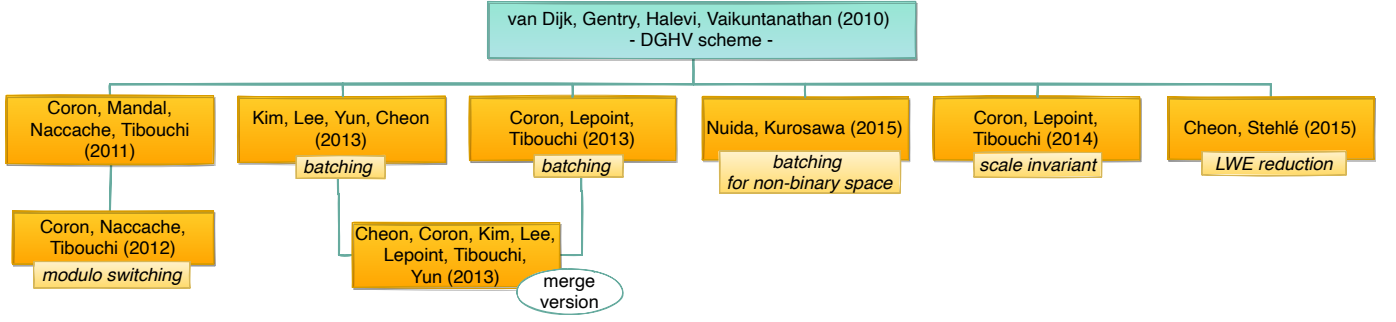


Fig. 5. Main FHE schemes based on the AGCD problem.

assumption with error-free problem. Informally, this problem is similar to the AGCD problem but with the stronger assumption that x_0 does not have the error r_0 considered in the conventional AGCD assumption. It is worth commenting that this is quantumly broken by Shor [68].

A *batch* version of DGHV scheme was independently proposed in 2013 by Kim, Lee, Yun and Cheon [69] and by Coron, Lepoint and Tibouchi [70] (the merged version is in [71]). Moreover, Nuida and Kurosawa [72] proposed the batching technique for non-binary message spaces. In 2014 Coron, Lepoint and Tibouchi [73] improved the DGHV scheme using the *scale-invariance* property given by Brakerski [47]. In 2015, Cheon and Stehlé [74], inspired by [47], introduced a reduction from LWE to AGCD and then presented a new AGCD-based fully homomorphic encryption scheme based on the hardness of this new variant. It is worth highlighting that this construction is the first DGHV variant that did not require the SSSP hardness assumption.

C. Second Generation: FHE based on LWE and RLWE

In 2011, Brakerski and Vaikuntanathan introduced, leveraging the bootstrapping technique, two FHE schemes based on the LWE [15] (extended version in 2014 [75]) and the RLWE [16] (see Sections III-G to III-I) problems, and also the concept of circular security. These works, described in Figure 6, initiated the second generation of FHE schemes. The LWE-based symmetric scheme described in [75], known as BV, is described as follows:

- *encryption*: the message $m \in \mathbb{F}_2$ is encoded into a ciphertext c such that

$$c = (a, b = \langle a, s \rangle + 2e + m) \in \mathbb{Z}_q^n \times \mathbb{Z}_q,$$

where e is the error randomly chosen from an error distribution χ and $s \in \mathbb{Z}_q^n$ is the secret key composed of random elements in \mathbb{Z}_q .

- *decryption*: outputs the plaintext $(b - \langle a, s \rangle \bmod q) \bmod 2$, which is equal to $(2e + m \bmod q) \bmod 2$. The decryption works properly if e is *small enough*, specifically $e < q/2$.

In this paper, the authors also introduced two novel techniques called *re-linearization* and *dimension-modulus reduction*. The re-linearization is needed to reduce the multiplication ciphertext size from almost $n^2/2$ back to regular size, i.e. $n + 1$. To obtain this reduction the authors transform the quadratic equation of $c_1 \cdot c_2$ into a linear equation by means of “encrypting” all the terms of the symmetric key under a new key. Later on, Brakerski, Gentry and Vaikunthanathan [11] called this technique *key switching*.

The *dimension-modulus reduction* technique (called also *modulus switching* [11] in subsequent works) converts an SHE into an FHE scheme transforming a ciphertext c modulo q into another ciphertext c' modulo p , where p is sufficiently smaller than q . Specifically, each element in \mathbb{Z}_q is converted into an element in \mathbb{Z}_p by first multiplying it by p/q , and then taking the closest integer. An interesting side effect of this operation is that the error in the ciphertext decreases³. The lower noise growth due to the adoption of modulus switching allows to homomorphically evaluate the decryption circuit without the squashing method proposed by Gentry. Thus, the SSSP assumption is no longer required (see Section IV-B).

In [16], Brakerski and Vaikunthanathan proposed a new version of the BV scheme, where the scheme security is based on the polynomial LWE (PLWE) problem. Note that the PLWE problem is equivalent to the RLWE problem [29] (see Section III-H). The main difference between the LWE version of the BV scheme and the RLWE version is that it represents the message, the ciphertext and the keys, as elements in R_q , where $R_q = \mathbb{Z}_q[x]/\langle f(x) \rangle$ with $f(x) \in \mathbb{Z}[x]$ is a polynomial of degree d and q is a prime.

In [11], Brakerski, Gentry and Vaikunthanathan proposed a method for defining a leveled fully homomorphic scheme (see Section IV-A), which avoids the computationally expensive bootstrapping technique. Based on this method, they defined the BGV scheme, and provided batching and modulus switching techniques. Additionally, the authors also gave a bootstrapping technique to transform the leveled version into an FHE version.

This new method made the scheme applicable to practical scenarios, thus fostering increasing interest from the research community. The authors introduced two variants of the BGV

³The error decreases when comparing ciphertexts before and after modulus switching, but the new ciphertext has also reduced modulus, and the error level relative to its modulus is actually higher after modulus switching since this technique introduces some error.

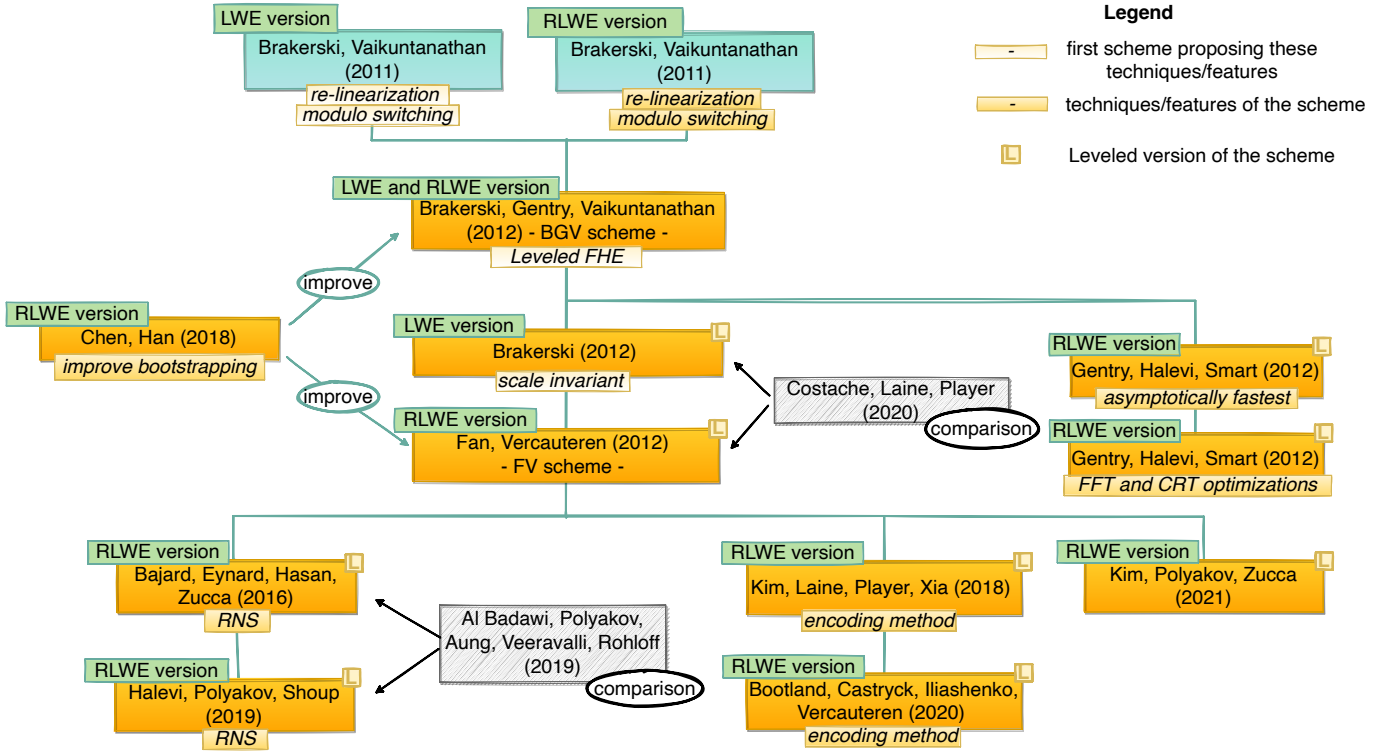


Fig. 6. Main second generation FHE schemes based on the LWE and RLWE problem.

scheme (one based on LWE and another on RLWE assumption). The RLWE-based version of BGV scheme, described in Scheme V.1, is more efficient than the LWE counterpart, and it is implemented in the most widely used FHE library HELib [76] (from IBM). The library, implemented by Shai Halevi and Victor Shoup, is open source and it enables the construction of a boolean circuit of any depth (more information about libraries can be found in Section V-G). It is worth commenting that the architecture of the library is based on a variant of the BGV scheme that introduces some optimizations, proposed by Gentry, Halevi and Smart [77], for the specific implementation of the AES circuit [78]. Also, in [79] Gentry, Halevi and Smart, using the batch techniques of Smart-Vercauteren [60] and Brakerski-Gentry-Vaikuntanathan [15], provided the asymptotically fastest scheme, with significant impact towards fast bootstrapping.

Scheme V.1: BGV RLWE-based scheme [11]

Let d be a power of 2, q be an odd positive integer modulus, and χ be an error distribution over R , where $R = \mathbb{Z}[x]/\langle x^d + 1 \rangle$. Let B be a bound (with overwhelming probability) on the length of elements outputted by χ . B is set to be as small as possible while maintaining security. For any natural integer p we write $R_p = \mathbb{Z}_p[x]/\langle x^d + 1 \rangle$. The scheme works as follows:

- *key generation*: takes as input the security parameter λ and randomly chooses a small

secret element $s \in \chi$ and sets the secret key $sk = s = (1, s) \in R_q^2$. It generates $a' \in R_q$ uniformly at random and computes $b = a's + 2e$ where e is a random error in χ . It outputs sk and $pk = a = (b, -a')$. Note that $\langle a, s \rangle = 2e$.

- *encryption*: takes as input the public key $pk \in R_q^2$ and the message $m \in R_2$. It converts m into a vector $\mathbf{m} = (m, 0) \in R_q^2$ and chooses random $r, e_0, e_1 \in \chi$. It outputs the ciphertext $\mathbf{c} = \mathbf{m} + 2(e_0, e_1) + ar$. Namely, $\mathbf{c} = (c_0, c_1) = (m + 2e_0 + br, 2e_1 - a'r) \in R_q^2$.
- *decryption*: takes as input the secret key $s \in R_q^2$ and the ciphertext $\mathbf{c} \in R_q^2$. It computes $\langle \mathbf{c}, s \rangle = c_0 + c_1 s = m + 2e_0 + 2e_1 s + 2er$ and outputs $((m + 2(e_0 + e_1 s + er)) \bmod q) \bmod 2 = m$. Note that the decryption works properly because e, e_0, e_1 and s are *small enough* (since are elements of χ).
- *(homomorphic) properties*.
 - The addition is a component-wise addition. The decryption works as long as the resulting error does not overlap the modulus q .
 - The multiplication is slightly more complicated. Let $L_y(x) : R_q^2 \rightarrow R_q$ be a linear map such that $L_y(x) = \langle y, x \rangle$. So, the

ciphertexts multiplication $\langle c_1, s \rangle \cdot \langle c_2, s \rangle = L_{c_1}(s)L_{c_2}(s)$ and the decryption of this multiplication is

$$\left(L_{c_1}(s)L_{c_2}(s) \right) \bmod q \bmod 2 = m_1 \cdot m_2.$$

$L_{c_1}(s)L_{c_2}(s)$ is a quadratic function in s where the c_i 's are the coefficients:

$$s^2 c_{1,1} c_{2,1} + s(c_{1,0} c_{2,1} + c_{1,1} c_{2,0}) + c_{1,0} c_{2,0},$$

which can be decrypted using a *long* secret key: $(s^2, s, 1)$. The inconvenience is that every multiplication expands the decryption key. So to reduce the decryption key, the authors use the key switching technique which consists in, given the keys s, s' and a ciphertext c , finding c' such that $\langle c, s \rangle = \langle c', s' \rangle \bmod q$ (for more details see [11, sec. 4.2]).

The leveled version of the schemes can be found in [11, sec. 4.4].

Brakerski [47] provided another variant of the BGV scheme, based on a technique called *scale invariant*. This technique reduces the error increase produced by homomorphic multiplications from exponential to linear. Intuitively, the idea behind the scale invariant technique is to scale down the ciphertext and the error by a factor of q , where q is the ciphertext modulus. This method replaces the modulus switching technique.

Another relevant contribution of [47] is a classical reduction to prove that the security of the scheme is based on the hardness of the GapSVP problem (see Section III-C). Previous schemes were, initially, proven secure only with quantum reductions (although now they also count with classical reductions).

The RLWE version of the Brakerski scheme was implemented and optimized by Fan and Vercauteren [12], and named FV scheme (see Scheme V.2). In the following sections of this work we refer to these two schemes as the B/FV scheme, namely for the LWE/RLWE variants. The FV scheme is one of the two schemes implemented in the Microsoft's Simple Encrypted Arithmetic Library (SEAL) [80], and it allows modular arithmetic to be performed on encrypted integers.

Following this research line, Bajard, Eynard, Hasan and Zucca [81] proposed an optimization, called RNS FV (BEHZ variant), when the ciphertext has large coefficients. It is based on the Residue Number System (RNS) since it uses CRT representation. This variant is improved in a subsequent work by Halevi, Polyakov and Shoup [82] (HPS variant). Both approaches were evaluated in [83] by Al Badawi, Polyakov, Aung, Veeravalli and Rohloff, where the authors show that the HPS variant [82] has a better decryption and homomorphic multiplication runtimes with respect to [81]. This improvement is specially significant in circuits with high multiplicative depth. Moreover, the multiplicative depth supported by the HPS variant is typically 15% – 30% larger than for BEHZ for the same parameters choice.

Chen and Han [84] improved the bootstrapping technique of both BGV and FV schemes for a large plaintext modulus (i.e. a large prime power). Another modification of the FV scheme is proposed in [85] by Kim, Laine, Player and Xia. The plaintext space is switched from R_t to \mathbb{Z}_{b^n+1} , by means of using the Hoffstein and Silverman trick [86] where the plaintext modulus t is substituted by a polynomial $x - b$, specifically:

$$R_t = \mathbb{Z}[x]/\langle x - b, x^n + 1 \rangle = \mathbb{Z}[x]/\langle x - b, b^n + 1 \rangle \cong \mathbb{Z}_{b^n+1}.$$

A recent generalization of this work is given in [87] where Bootland, Castryck, Iliashenko and Vercauteren proposed a plaintext modulus as $x^m + b$ instead of $x - b$. Both works, [85] and [87], follow a strand of articles ([88]–[90]) studying an encoding method for transforming a real input data (namely, integer, rational or a complex number) into a polynomial, which is an element of the message space (i.e. plaintext) of a RLWE scheme. It is worth highlighting that both [85] and [87] achieve a reduction of the error growth when compared to the original version of the FV scheme and, as a consequence, both of them enable the evaluation of circuits with higher multiplicative depth. The main difference between these two approaches is that whereas [85] encodes fractional numbers, [87] encodes complex numbers.

Latterly, Kim, Polyakov and Zucca [91] proposed several optimizations to the FV and the BGV schemes and suggest a different approach to compute the ciphertext modulus of the BGV scheme, which does not require dynamic noise estimation. In this setting, their FV variant has better noise growth than BGV for all plaintext moduli but their FV variant is faster than BGV only for small plaintexts, while BGV is faster for intermediate and large plaintexts. This latter result confirms the previous work of Costache, Laine and Player [92] where the authors, extending a previous performance evaluation study of Costache and Smart [93], compared the error growth for both schemes. Specifically, in [92] the authors proved that, for a particular small plaintext modulus and circuit depth, BGV required a larger parameter set than FV. Namely, BGV outperforms FV when the plaintext modulus is above some value t , for $t \in [2^{32}, 2^{64}]$. However, this analysis does not consider some of the available optimizations for BGV and FV, and hence further performance evaluations should be considered in the future.

Scheme V.2: FV scheme [12]

Let $R = \mathbb{Z}[x]/\langle x^d + 1 \rangle$ where d is a power of 2. Let q and p be positive integers, let $\Delta = \lfloor q/p \rfloor$ and $r_t(q) = q \bmod p$. For any natural number t we write $R_t = \mathbb{Z}_t[x]/\langle x^d + 1 \rangle$. Let χ be a B -bounded probability distribution over R_q . Then, the FV scheme is constructed as follows:

- *key generation*: takes as input the security parameter λ and outputs the small secret key $sk = s \in \chi$. It generates $a \in R_q$ uniformly at

random and computes $-(a \cdot s + e) \bmod q$ where $e \in \chi$ is a small random error. It outputs sk and $pk = (p_0, p_1) = (-(a \cdot s + e) \bmod q, a)$.

- *encryption*: takes as input the message $m \in R_p$ and the public key $pk \in R_q^2$. It chooses at random the values $u, e_1, e_2 \in \chi$ and outputs the ciphertext $c = (c_0, c_1)$, where $c_0 = (p_0 \cdot u + e_1 + \Delta m) \bmod q$ and $c_1 = (p_1 \cdot u + e_2) \bmod q$.
- *decryption*: takes as input the secret key $s \in R_q$ and the ciphertext $c \in R_q^2$. It computes $\lfloor \frac{p \cdot (c_0 + c_1 \cdot s) \bmod q}{q} \rfloor \bmod p = m$.
- *(homomorphic) properties*. An extensive description can be found in [12, sec. 4]), here we cater for a brief overview. The ciphertext c can be seen as a polynomial evaluated in s , i.e. $c(s)$, instead of a vector with two components.
 - The addition is trivial: $c_1(s) + c_2(s) \bmod q$
 - The multiplication of two ciphertexts gives as a result a quadratic polynomial

$$c_1(s) \cdot c_2(s) = \alpha_0 + \alpha_1 \cdot s + \alpha_2 \cdot s^2$$

which can be transformed into a decryptable ciphertext by means of a *re-linearization* process, which reduces by one the degree of the ciphertext.

which supports computation on ciphertexts encrypted under different keys (for more details on Multi-Key FHE see [24]). This scheme, called LTV, uses the bootstrapping and modulus switching techniques and it is constructed as follows:

- *key generation*: chooses two small random polynomials $f', g \in \chi$, where χ is a B-bounded distribution over $R = \mathbb{Z}[x]/\langle x^d + 1 \rangle$ and d is a power of 2. It outputs the secret key $f = 2f' + 1 \in R$, where $f \equiv 1 \bmod 2$ and f is invertible in R_q , and the public key $h = 2gf^{-1} \bmod q \in R_q$.
- *encryption*: computes the ciphertext as $c = hs + 2e + m \in R_q$, where $m \in \mathbb{F}_2$ is the message and s and e are random small elements in χ . The ciphertext is an element of $R_q = \mathbb{Z}_q[x]/\langle x^d + 1 \rangle$.
- *decryption*: computes $m = (fc \bmod q) \bmod 2$.

The security of this scheme is based on circular security, the RLWE problem and the *Decisional Small Polynomial Ratio* (DSPR) problem. The DSPR problem states that it is hard to distinguish between h (as defined in the scheme construction) and an uniformly random polynomials in R_q . One year later Bos, Laute, Loftus and Naehrig [98] modified the LTV scheme [97] proposing two schemes: i) YASHE (i.e. Yet Another Somewhat Homomorphic Encryption scheme), for which they removed the DSPR assumption using the scale invariant at the cost of having a large evaluation key and a complex key switching method; ii) a YASHE version including again the DSPR assumption to achieve a more practical construction. In 2016, Albrecht, Bai, and Ducas [99] and, in an independent work, Cheon, Jeong and Lee [100], provided a *sub-field lattice* attack which renders any NTRU-like scheme based on the DSPR problem insecure, for some particular parameters choice. Finally, Doröz and Sunar [101], adapted the GSW scheme [54] (see Section V-E) to the NTRU setting by removing the DSPR assumption.

Although NTRU-like schemes are, in general, faster than RLWE schemes, the work by Lepoint and Naehrig [102] showed that the FV scheme [17] has lower error than the YASHE scheme [98]. Also, Kim and Lauter [103] proved that YASHE is better than BGV for a low-degree computation but BGV is more efficient than YASHE for high circuit depths. Moreover, Costache and Smart [93] demonstrated that YASHE is more efficient than BGV for small plaintexts modulus, but BGV is more efficient for large plaintexts. Figure 7 shows the diagram of second generation schemes based on NTRU, including the most relevant papers providing performance comparisons. To conclude this section, it is worth commenting that the parameters choice for NTRU-based schemes is affected by the attacks proposed in [99], [100]. Namely, to obtain a secure NTRU-based scheme the parameters should be significantly increased with respect to the sizes proposed before these attacks were published. This made this kind of schemes less efficient.

D. Second Generation: FHE based on NTRU

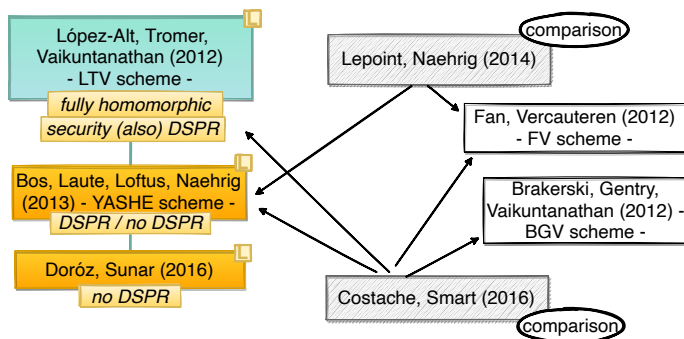


Fig. 7. Main FHE schemes based on NTRU.

NTRU [94] is a lattice-based encryption scheme introduced by Hoffstein, Pipher and Silverman in 1996, with a provisional patent filed, and granted in 2000 [95]. Since its inception, the security of this scheme was under discussion by the research community, until 2011, when Stehlé and Steinfeld [96] slightly modified the scheme to obtain a variant which security is based on the RLWE assumption. One year later, López-Alt, Tromer and Vaikuntanathan [97] introduced the first FHE scheme inspired from the Stehlé-Steinfeld NTRU variant. It is considered part of the second generation of FHE schemes. Specifically, the authors propose a new notion of homomorphic encryption scheme called *Multi-Key FHE*,

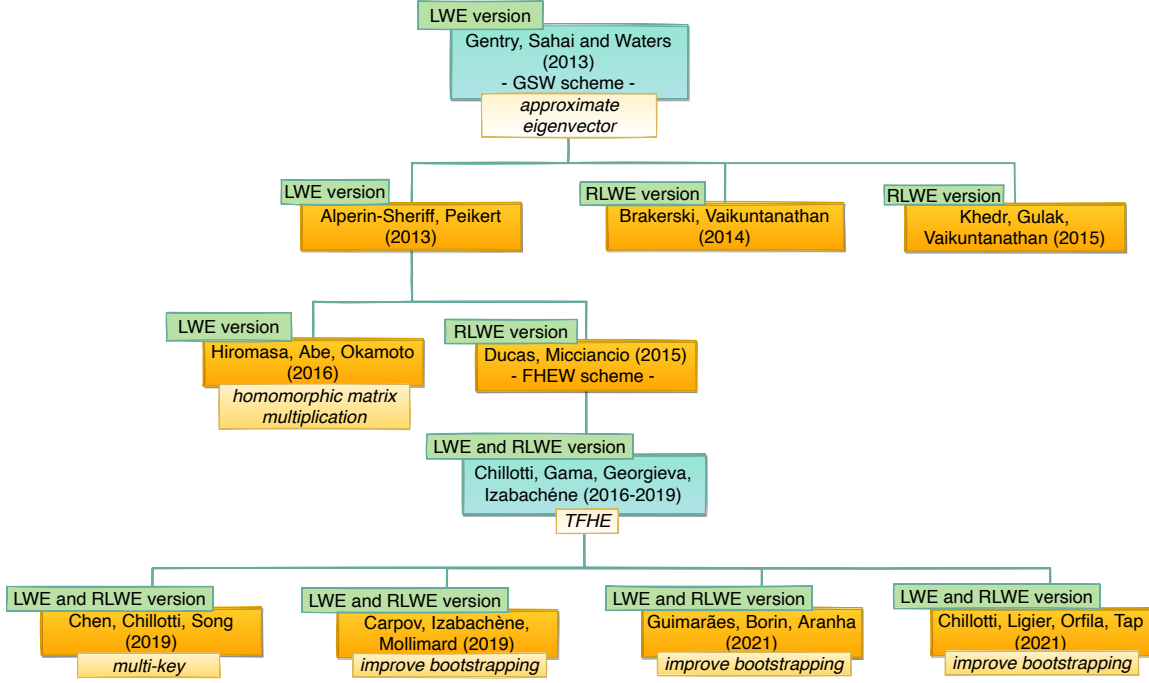


Fig. 8. Main third generation of FHE schemes based on the LWE and RLWE problems.

E. Third Generation: FHE based on LWE and RLWE

A second family of (R)LWE schemes (also called third generation FHE, as Peikert notes in [28]) started with the Gentry, Sahai and Waters [54] (GSW) scheme. The GSW scheme proposes a different approach to perform homomorphic operations, introducing the *approximate eigenvector method*, which removes the requirement for key and modulus switching techniques. This new technique reduces the error growth introduced by homomorphic multiplications to a small polynomial-factor. As shown in [104], when multiplying ℓ ciphertexts, all starting with the same error level, the final error grows by a ℓ poly(n) factor, where n is the dimension of the scheme (i.e. the dimension of the lattice). This is a distinctive aspect with respect to previous schemes, such as BGV or FV, for which the final error grows by a quasi-polynomial-factor. The RLWE version of this scheme was given by Khedr, Gulak and Vaikuntanathan in [105]. Figure 8 describes this second family of (R)LWE schemes.

Adopting the notations used in [106] and [57], the (simplified) GSW scheme construction⁴ is as follows:

- *key generation*: outputs the secret key $\mathbf{s} = (1, s_2, \dots, s_n) \in \mathbb{Z}_q^n$, where the s_i 's are chosen at random and a public key $A \in \mathbb{Z}_q^{n \times n}$ that is a matrix $n \times n$ such that $A \cdot \mathbf{s} = \mathbf{e} \approx 0$;
- *encryption*: computes the ciphertext $C = mI_n + RA$, where $m \in \mathbb{Z}_q$ is the message, I_n is the identity matrix, R is a random matrix with size $n \times n$ and with coefficients in \mathbb{F}_2 . This means that the entries of R are

small;

- *decryption*: First, computes $C\mathbf{s} = mI_n\mathbf{s} + RA\mathbf{s} = mI_n\mathbf{s} + R\mathbf{e} \approx mI_n\mathbf{s}$. Note that because R is small, if $A\mathbf{s} \approx 0$ then $RA\mathbf{s} \approx 0$. Finally, outputs the first element of the vector $x \approx mI_n\mathbf{s} \approx (ms_1, \dots, ms_n)$, which is m , since $s_1 = 1$.

The main drawbacks of the GSW scheme are the high communication costs (the ciphertext is large with respect to the corresponding plaintext) and the computation complexity. To reduce the computational overhead, various optimizations have been proposed to improve the bootstrapping procedure (see Figure 8). Specifically, Alperin-Sheriff and Peikert [106], [107] suggested a new bootstrapping algorithm considering decryption as an *arithmetic* function instead of a boolean circuit. Hiromasa, Abe and Okamoto [108] optimized the Alperin-Sheriff and Peikert procedure [106] constructing an FHE scheme (based on GSW scheme) that supports homomorphic matrix operations. Moreover, Brakerski and Vaikunthanathan [104], starting from GSW, proposed the first work that managed to get FHE scheme based on GapSVP with polynomial approximation factors (and circular security). Ducas and Micciancio [109] proposed a ring variant, the FHEW scheme, of the Alperin-Sheriff and Peikert bootstrapping technique [106] and introduced a new method to homomorphically compute the NAND of two standard LWE ciphertexts (with standard we refer to the ones of Regev's scheme [41]). In this work, they also adopt the complex FFT which enables the implementation of the scheme with the *Fastest Fourier Transform in the West* [110] library (the “W” of the scheme's name FHEW comes from this). These set of optimizations render the GSW's scheme bootstrapping procedure faster than the BGV's scheme

⁴Note that, for the sake of clarity, this version is a simplified description of the scheme and it is not homomorphic for multiplication. To get homomorphic multiplications we would also need the bit decomposition.

one.

In a subsequent paper, Chillotti, Gama, Georgieva and Izabachène [13] improved the Ducas-Micciancio's result. The authors proposed three different FHE schemes over the Torus, generally called TFHE: i) TLWE, which is a generalized version of the LWE problem for the Torus; TRLWE, which is its ring variant; and iii) TRGSW, which improves the ring version of GSW scheme. The messages of the TLWE scheme are in the torus \mathbb{T} and the ciphertexts in \mathbb{T}^{n+1} , whereas the ring version TRLWE of this scheme works with plaintexts in the R -module, where R is the ring of integer polynomials. Namely, the message space is $T = \mathbb{R}[x]/\langle x^d + 1 \rangle \bmod 1$. The TRGSW scheme encrypts elements of the ring of integer polynomials into a vector of TRLWE ciphertexts, namely, $C \in T^{(k+1)\ell}$. It is worth commenting that [13] is an extended and improved version of previous articles presented in Asiacrypt 2016 [111], where they speed up the bootstrapping procedure and reduce the bootstrapping key size compared to FHEW; and Asiacrypt 2017 [112], where they improve the leveled version of TFHE.

A detailed description of the TRLWE scheme is provided in Scheme V.3.

Scheme V.3: TRLWE scheme [13]

Let us consider the following notations: $R = \mathbb{Z}[x]/\langle x^d + 1 \rangle$, where d is a power of 2, $T = \mathbb{R}[x]/\langle x^d + 1 \rangle \bmod 1$ and $R_2 = \mathbb{F}_2[x]/\langle x^d + 1 \rangle$, that is, any element in R_2 is a polynomial in R with binary coefficients (in a recent work [113], the authors used \mathbb{Z}_q instead of a torus). Then, the TRLWE scheme is constructed as follows:

- *key generation*: takes as input the security parameter λ and outputs the small secret key $\mathbf{s} \in R_2^n$.
- *encryption*: takes as input the secret key $\mathbf{s} \in R_2^n$, the error parameter α , and the message $m \in T$. Then, it chooses a uniformly random mask $\mathbf{a} \in T^n$, and a small error $e \in \chi$, where χ is a B-bounded distribution. Then it outputs the ciphertext.

$$\mathbf{c} := (\mathbf{a}, \mathbf{s} \cdot \mathbf{a} + m + e) \in T^n \times T.$$

- *decryption*: takes as input the secret key $\mathbf{s} \in R_2^n$ and the ciphertext $\mathbf{c} \in T^{n+1}$. Then it computes the secret linear κ -Lipschitz function $\varphi_{\mathbf{s}}$ (called *phase*) of the ciphertext \mathbf{c} . The phase $\varphi_{\mathbf{s}} : T^n \times T \rightarrow T$ is such that $\varphi_{\mathbf{s}}(\mathbf{a}, b) = b - \mathbf{s} \cdot \mathbf{a}$. Note that this function is parametrized by the secret key $\mathbf{s} \in R_2^n$. The phase $\varphi_{\mathbf{s}}(\mathbf{c})$ is close to the actual message:

$$\varphi_{\mathbf{s}}(\mathbf{c}) = \mathbf{s} \cdot \mathbf{a} + m + e - \mathbf{s} \cdot \mathbf{a} = m + e.$$

To conclude, it rounds $\varphi_{\mathbf{s}}(\mathbf{c})$ to the nearest point in the message space $M \subset T$.

- *(homomorphic) linear combinations of ciphertexts*. Let $\mathbf{c}_1, \dots, \mathbf{c}_p$ be p independent ciphertexts under the same key \mathbf{s} , and let f_1, \dots, f_p be integer polynomials in R .

We consider $\mathbf{c} = \sum_{i=1}^p f_i \cdot \mathbf{c}_i$ such that the error amplitude remains smaller than $1/4$, that is, $\|e\|_{\infty} \leq 1/4$. Then by [13, Fact 3.5], \mathbf{c} is a ciphertext and

$$\text{Dec}_{\mathbf{s}}(\mathbf{c}) = \sum_{i=1}^p f_i \cdot \text{Dec}_{\mathbf{s}}(\mathbf{c}_i).$$

As the authors of [13] pointed out, the ciphertexts can be linearly combined to obtain a new ciphertext which is the linear combination of the messages. But when we have to manipulate the ciphertext non-linearly, TLWE seems to miss some properties. In order to avoid this problem, the authors of [13] proposed the generalized scale invariant version of GSW, called TRGSW. Note that to “switch” from TRLWE to TLWE (and vice-versa) we only have to consider the real torus instead of T , $\{0, 1\}$ instead of B and \mathbb{Z} instead of R .

A relevant feature of the TFHE scheme is that the bootstrapping technique enables an univariate function to be evaluated simultaneously to the noise reduction operation [113]. This is normally referred to as programmable bootstrapping (PBS). Several optimizations have been proposed to improve the TFHE scheme and, in particular, its programmable bootstrapping procedure. Namely, in [114], Carpv, Izabachène and Mollimard proposed a multi-output version of the PBS, that is, a new technique to perform several homomorphic operations over different variables with a single bootstrapping execution. This construction can also be used to evaluate homomorphically several functions over the same encrypted message simultaneously. In a recent work, Guimarães, Borin and Aranha [115] optimized the bootstrapping procedure to evaluate multiple functions on large ciphertexts and Chillotti, Ligier, Orfila and Tap [113] proposed several enhancements, including a generalized method to evaluate several functions at once without adding additional error. Finally, in [116], Chen, Chillotti and Song proposed a multi-key homomorphic encryption scheme from TFHE. They provide two methods to multiply a ciphertext encrypted with a single key by a ciphertext encrypted with multiple keys. To conclude this section, we would like to refer interested readers and TFHE practitioners to the recently published guide of Joye [117], which presents implementation details, theoretical examples and a clear description of the programmable bootstrapping technique.

F. Fourth Generation: FHE based on LWE and RLWE

In 2017 a new generation of FHE schemes (see Figure 9) was introduced by Cheon, Kim, Kim and Song [14]. In this work, the authors proposed a method to construct a leveled Homomorphic Encryption scheme for Approximate Arithmetic Numbers (HEAAN, but also called CKKS scheme from the

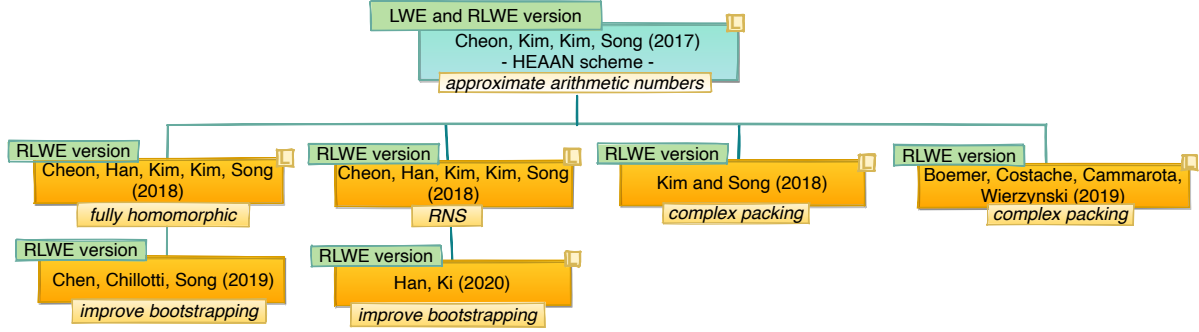


Fig. 9. Main fourth generation FHE schemes based on the LWE and RLWE problems.

names of authors). The scheme construction is described in Scheme V.4. The scheme was extended one year later to a fully homomorphic encryption scheme by Cheon, Han, Kim, Kim and Song [118], and, subsequently [119] the same authors presented a variant of HEAAN scheme by means of including a ciphertext packing technique based on the CRT. Also, the bootstrapping version [118] of HEAAN was enhanced in [120] by Chen, Chillotti and Song, whereas, in [121], Han and Ki discussed and improved the bootstrapping version of [119]. In [122] Boemer, Costache, Cammarota and Wierzynski introduced several optimizations, based on complex packing, to improve the runtime of scalar encoding and of ciphertext-plaintext addition and multiplication operations. Moreover, a different kind of complex packing was introduced by Kim and Song in [123].

In a recent work, Li and Micciancio [124] presented an attack against HEAAN that works for a specific application scenario. The authors proved that the HEAAN scheme (as well as all improvements and implementations of approximate encryption schemes) is vulnerable to an attack by an adversary that has access to the encryption functionality. Namely, the adversary can extract a secret key using only linear algebra or lattice reduction techniques. Specifically, the secret key can be obtained if the decrypted element and corresponding ciphertext are both known, because the error is a linear combination of the secret key's elements (see Scheme V.4). However, as Cheon, Hong and Kim pointed out in [125], this attack can be prevented if the owner of the secret key does not share the result of the decrypted messages. Unfortunately, some applications such as secure multi-party computation or differential privacy techniques require sharing some plaintexts. In such cases, Li and Micciancio suggested that their attack could be probably avoided by modifying the decryption function by means of adding an error at the end of the decryption process, as suggested in [124].

Scheme V.4: HEAAN scheme [14]

Let $R = \mathbb{Z}[x]/\langle x^d + 1 \rangle$ and $d = 2^M$. For a base p , a modulus q_0 and a natural integer L (chosen level), let $q_\ell = p^\ell \cdot q_0$ for $\ell = 1, \dots, L$. Note that a ciphertext of level ℓ is a vector in R_{q_ℓ} . Let us consider the following relevant distributions. For a real number σ , $DG(\sigma^2)$

is a vectorial discrete Gaussian distribution over \mathbb{Z}^d , which samples each of its components from independent discrete Gaussian distributions of variance σ^2 . For a real number $0 < \rho < 1$, the distribution $ZO(\rho)$ is the distribution over $\{-1, 0, 1\}^d$ which draws 0 with probability $1 - \rho$ and -1 or 1 with probability $\rho/2$. Finally, let us consider χ a B -bounded distribution.

The leveled HEAAN scheme is constructed as follows:

- *key generation*: takes as input the security parameter λ and chooses M , integers h and t , and a real number σ such that the RLWE problem achieves a complexity of 2^λ . It computes a secret key $\text{sk} = (1, s)$, where $s \in \chi$. It generates $a \in R_{q_L}$ uniformly at random and computes $-as + e \bmod q_L$, where e is $DG(\sigma^2)$. Finally, it samples $a' \in R_{t \cdot q_L}$ and $e' \in DG(\sigma^2)$ and computes $b' = -as + e' + ts' \bmod t \cdot q_L$. It outputs the secret key $\text{sk} = (1, s)$, the public key $\text{pk} = (b, a)$ and the evaluation key $\text{evk} = (b', a')$.

- *encryption*: takes as input the message $m \in R$ and the public key $\text{pk} \in R_{q_L}^2$. It chooses at random the values $v \in ZO(1/2)$ and $e_0, e_1 \in DG(\sigma^2)$ and outputs

$$c = (\beta, \alpha) = v\text{pk} + (m + e_0, e_1) \bmod q_\ell.$$

- *decryption*: takes as input the secret key $\text{sk} = (1, s)$ and the ciphertext $\mathbf{c} \in R_{q_\ell}^2$. It computes $m = \langle \mathbf{c}, \text{sk} \rangle \bmod q_\ell = \beta + \alpha s \bmod q_\ell$.

- *(homomorphic) properties*.

- The addition is trivial: $\mathbf{c}_1 + \mathbf{c}_2$
- The multiplication of two ciphertexts $\mathbf{c}_i = (\beta_i, \alpha_i)$ with $i = 1, 2$ is

$$\mathbf{c}_1 \cdot \mathbf{c}_2 = (d_0, d_1) + \lfloor t^{-1} d_2 \text{evk} \rfloor \bmod q_\ell,$$

$$\text{where } (d_0, d_1, d_2) = (\beta_1 \beta_2, \alpha_1 \beta_2 + \alpha_2 \beta_1, \alpha_1 \alpha_2) \bmod q_\ell.$$

It is worth clarifying that when we have two ciphertexts \mathbf{c}, \mathbf{c}' of two different levels $\ell' < \ell$, we should reduce the level of the ciphertext with

a higher level to match both levels, i.e. $\ell' = \ell$. This can be achieved with a *rescaling* procedure that takes a ciphertext $c \in R_{q_\ell}^2$ at level ℓ and outputs $c' = \lfloor q_{\ell'}/q_\ell \rfloor \bmod q_{\ell'}$.

An interesting feature of HEAAN is that the message space can be represented as elements in the extension field \mathbb{C} . Informally, the message m can be embedded in $S = \mathbb{R}[x]/\langle x^d+1 \rangle$. Since, the roots of x^d+1 are the complex primitive roots of unity in \mathbb{C} , to convert the $m \in S$ into a vector of complex numbers it is sufficient to evaluate it at these complex roots. For more details see [14, sec. 3].

The fourth generation schemes are similar to the second generation. The main difference is that the fourth generation schemes are approximate schemes, i.e. they use approximate computation, which is considerably faster. Specifically, the fourth generation embeds the message space into a complex hyperplane, and the error during encryption is inserted as part of the approximation error that is inherently introduced during a computation over real-valued numbers. An interesting feature of HEAAN is the capability to homomorphically operate over approximations of real numbers, which makes it a suitable scheme to work with floating-point arithmetic. Another similarity is that the schemes from both generations have efficient packing techniques, and they can only compute fast sum and product (any non-linear operation becomes computationally expensive).

It is important to highlight the work proposed by Boura, Gama, Georgieva and Jethchev [126], the CHIMERA scheme, which is a hybrid solution combining three RLWE-based FHE schemes: TFHE [13], B/FV [12], [47] and HEAAN [118]. CHIMERA has the special property that it enables the switching between the three schemes. The authors start by defining a common plaintext space between the three schemes by constructing an embedding⁵ of the different message spaces. By leveraging the bootstrapping technique, CHIMERA enables switching ciphertexts from TFHE to FV (and vice-versa) and from HEAAN to FV (and vice-versa). FV must be used as intermediate step for transformations between TFHE and HEAAN. CHIMERA was first presented as a solution to the Idash'18 Track 2 competition [127], [114].

G. Final Considerations

To the best of the authors knowledge, BGV [11], B/FV [12], TFHE [13] and HEAAN [118] are the most practical and widely adopted schemes. The second generation schemes, BGV and B/FV, are suitable to work with finite fields in the modular exact arithmetic. They are equipped with efficient packing which enables the use of SIMD (namely, single instruction multiple data) instructions to perform computations over vectors of integers (i.e. batching). Thus, these schemes

are excellent candidates when large arrays of numbers are to be processed simultaneously.

Second generation schemes are not good candidates for circuits where bootstrapping is required (i.e. circuits with large multiplicative depth), or where non-linear functions are to be implemented. Third generation schemes should be adopted instead, namely TFHE, which can outperform previous schemes for bit-wise operations, i.e. when computations are expressed as boolean circuits [55]. The main limitation of TFHE is the lack of support for CRT packing (i.e. batching), hence the scheme can be outperformed by previous approaches when processing large amounts of data simultaneously. The fourth generation, i.e. HEAAN, is the best option for real numbers arithmetic. Table I provides a comparison among the schemes' families, and Figure 10 depicts the main applications for each generation of schemes.

SCHEMES	2nd Generation		3rd Generation	4th Generation
	BGV	B/FV	TFHE	HEAAN
PROS / APPLICATIONS	Integer Arithmetic		Bitwise operations	Real Number Arithmetic
	efficient packing (SIMD)		efficient boolean circuits	fast polynomial approx.
	fast escalar multiplication		fast bootstrapping	fast multiplicative inverse
	fast linear functions		fast number comparison	efficient DFT
	efficient leveled design			efficient logistic regression
CONS	slow bootstrapping		no support for batching	slow bootstrapping
	slow non-linear functions			slow non-linear functions

Fig. 10. Pros/cons of FHE schemes by generation.

H. Other Works

In a recent work by Doröz, Hoffstein, Pipher, Silverman, Sunar, White and Zhang [128], an FHE scheme based on a new hard problem was introduced: *Finite Field Isomorphism Problems*, that is based on the difficulty of recovering a secret isomorphism between two finite fields. Moreover, in 2019 Joux [129] proposed a scheme whose techniques are similar to those of (R)LWE, but with arithmetic modulo large *Fermat numbers*, namely numbers given by the expression $F = 2^{2^f} + 1$, where $f \in \mathbb{N}$.

VI. SECURITY OF SCHEMES BASED ON (R)LWE PROBLEMS

As described in Section III-G, the LWE problem consists of finding the secret vector $s \in \mathbb{Z}_q^n$, given $b \in \mathbb{Z}_q^m$ and $A \in (\mathbb{Z}_q)^{m \times n}$ such that $As + e = b \bmod q$, where $e \in \mathbb{Z}_q^m$ is sampled from the error distribution χ . The security of LWE-based schemes depends on the intractability of this problem, and attacks on these schemes are based on finding efficient algorithms to solve them. In this framework, Albrecht, Player and Scott [130], presented three different methodologies to solve the LWE-problem (Figure 11): i) based on the BDD problem (Section VI-B); ii) based on the SIS problem (Section VI-C); and iii) a direct search of the secret

⁵An instance of a mathematical structure that is contained within another instance

TABLE I
PROPERTIES OF THE MOST WIDELY ADOPTED FHE SCHEMES

Scheme	Fast operations			Fast packing/batching	Leveled design	Fast bootstrapping
	scalar mult	arithmetic	non arithmetic			
Second Generation (e.g. BGV, B/FV)	●	●	○	●	●	○
Third Generation (e.g. FHEW, TFHE)	●	●	●	○	●	●
Fourth Generation (e.g. HEAAN)	●	●	○	●	●	○

s (Section VI-D). Conceptually, the central part of the first and the second methodology is based on a lattice reduction. Namely, starting from a bad (i.e. long) lattice basis, find a better (i.e. reduced and more orthogonal) basis. Note that, in [130], the authors showed that there is no single-best attack against all possible parameters.

Regarding the schemes based on the RLWE problem, the same considerations apply. According to the Homomorphic Encryption Security Standard [55], if we choose correctly the error distribution then there are no better attacks on RLWE than on LWE. This is because the best known attacks do not leverage any property of the ring structure. In this claim, the correct choice of the error distribution only refers to a sufficiently *well spread* distribution [131]–[133] (Section VI-E).




Fig. 11. Different solving approach to LWE problem

A comprehensive explanation of lattice attacks can be found in [130], [134]–[136]. Moreover, the work of Bindel, Buchmann, Göpfert and Schmidt [137] extends previous studies with an analysis of how the number of samples affects the hardness of LWE. It is worth highlighting the work of Albrecht, Player and Scott [130], which not only describes in detail LWE attacks, but also provides a software tool to determine the security level of LWE instances. The first version of this tool is referred to as the LWE Estimator⁶, and has been adopted in the Homomorphic Encryption Security Standard [55] to provide specific parameters for FHE schemes. The new version, denoted by Lattice Estimator⁷, “was born out of frustration with the limitations of the old codebase” as Albrecht mentioned in his blog [138]. Figure 12 shows the timeline and fundamentals of the main lattice attacks proposed until the present time.

A. Lattice Reduction Algorithms

The most well-known lattice reduction algorithm used in practice is BKZ (block Korkin-Zolotarev reduction) due to Schnorr and Euchner [139]. This is a block-wise iterative algorithm for basis reduction that can be seen as a generalization of

the LLL algorithm introduced in 1982 by Lenstra, Lenstra and Lovász [140]. Recently, several variants of the BKZ algorithm have been proposed [141]–[145]. In these algorithms, the time complexity and the outcome quality (i.e. the orthogonality of the reduced basis) is characterised by the Hermite factor [146] and is given as a trade-off. Specifically, the run time of the BKZ algorithm is higher when the root-Hermite factor δ_0 is smaller [139]. This is also shown in Lindner and Peikert’s estimation [147] (Figure 12). This result is also supported by a more realistic estimation⁸ provided in [130]. Starting from the data provided by Liu and Nguyen [148], a similar relation between δ_0 and the run time of the BKZ 2.0 algorithm [141] was found by Albrecht, Cid, Faugère, Fitzpatrick and Perret [149] (Figure 12). It is also worth commenting that the quality of the output decreases with higher values of δ_0 . The formula linking the root-Hermite factor attainable by BKZ and the block-size b of this algorithm was heuristic established (and well-verified experimentally) by Chen in his PhD thesis [150]:

$$\delta_0 = \left((\pi b)^{\frac{1}{b}} \cdot \frac{b}{2\pi e} \right)^{\frac{1}{2(b-1)}}.$$

For more details about lattice reduction algorithms we refer the readers to [130] by Albrecht, Player and Scott, and the survey on algorithms for the SVP and CVP by Hanrot, Pujol and Stehlé [151].

B. Attacks Based on Bounded Distance Decoding Problem

These attacks are based on solving LWE by solving the BDD problem (Section III-D). Specifically, the main strategy of this kind of attacks consists in finding \mathbf{v} , the closest vector to $A\mathbf{s} + \mathbf{e}$, for a lattice $\mathcal{L} = \mathcal{L}(A)$. Note that knowing \mathbf{v} , which equals $A\mathbf{s}$, we can obtain \mathbf{s} and thus solve the LWE problem. The following strategies have been proposed:

1) *BDD enumeration (decoding)*: This attack was proposed by Lindner and Peikert [147], who modified the *Nearest Plane* algorithm analyzed by Babai [152], using multiple planes to decrease the failure probability of finding the vector $A^t \mathbf{s}$. The Lindner-Peikert algorithm has two main steps. First, it applies lattice reduction (using BKZ) to obtain a new basis $\{\beta_1, \dots, \beta_n\}$. Second, it performs a recursive search for an integer combination of the basis vectors β_i that is close to the target vector \mathbf{v} . In 2013, Liu and Nguyen [148], starting from the paper by Gama, Nguyen and Regev [153], ameliorated the Lindner-Peikert algorithm as well as Babai’s.

⁸As Albrecht pointed out [134] the LP model for estimating the cost of lattice-reduction is not correct for several reasons and the formula proposed by Lindner and Peikert turns out to be too optimistic.

⁶<https://bitbucket.org/malb/lwe-estimator>

⁷<https://github.com/malb/lattice-estimator>

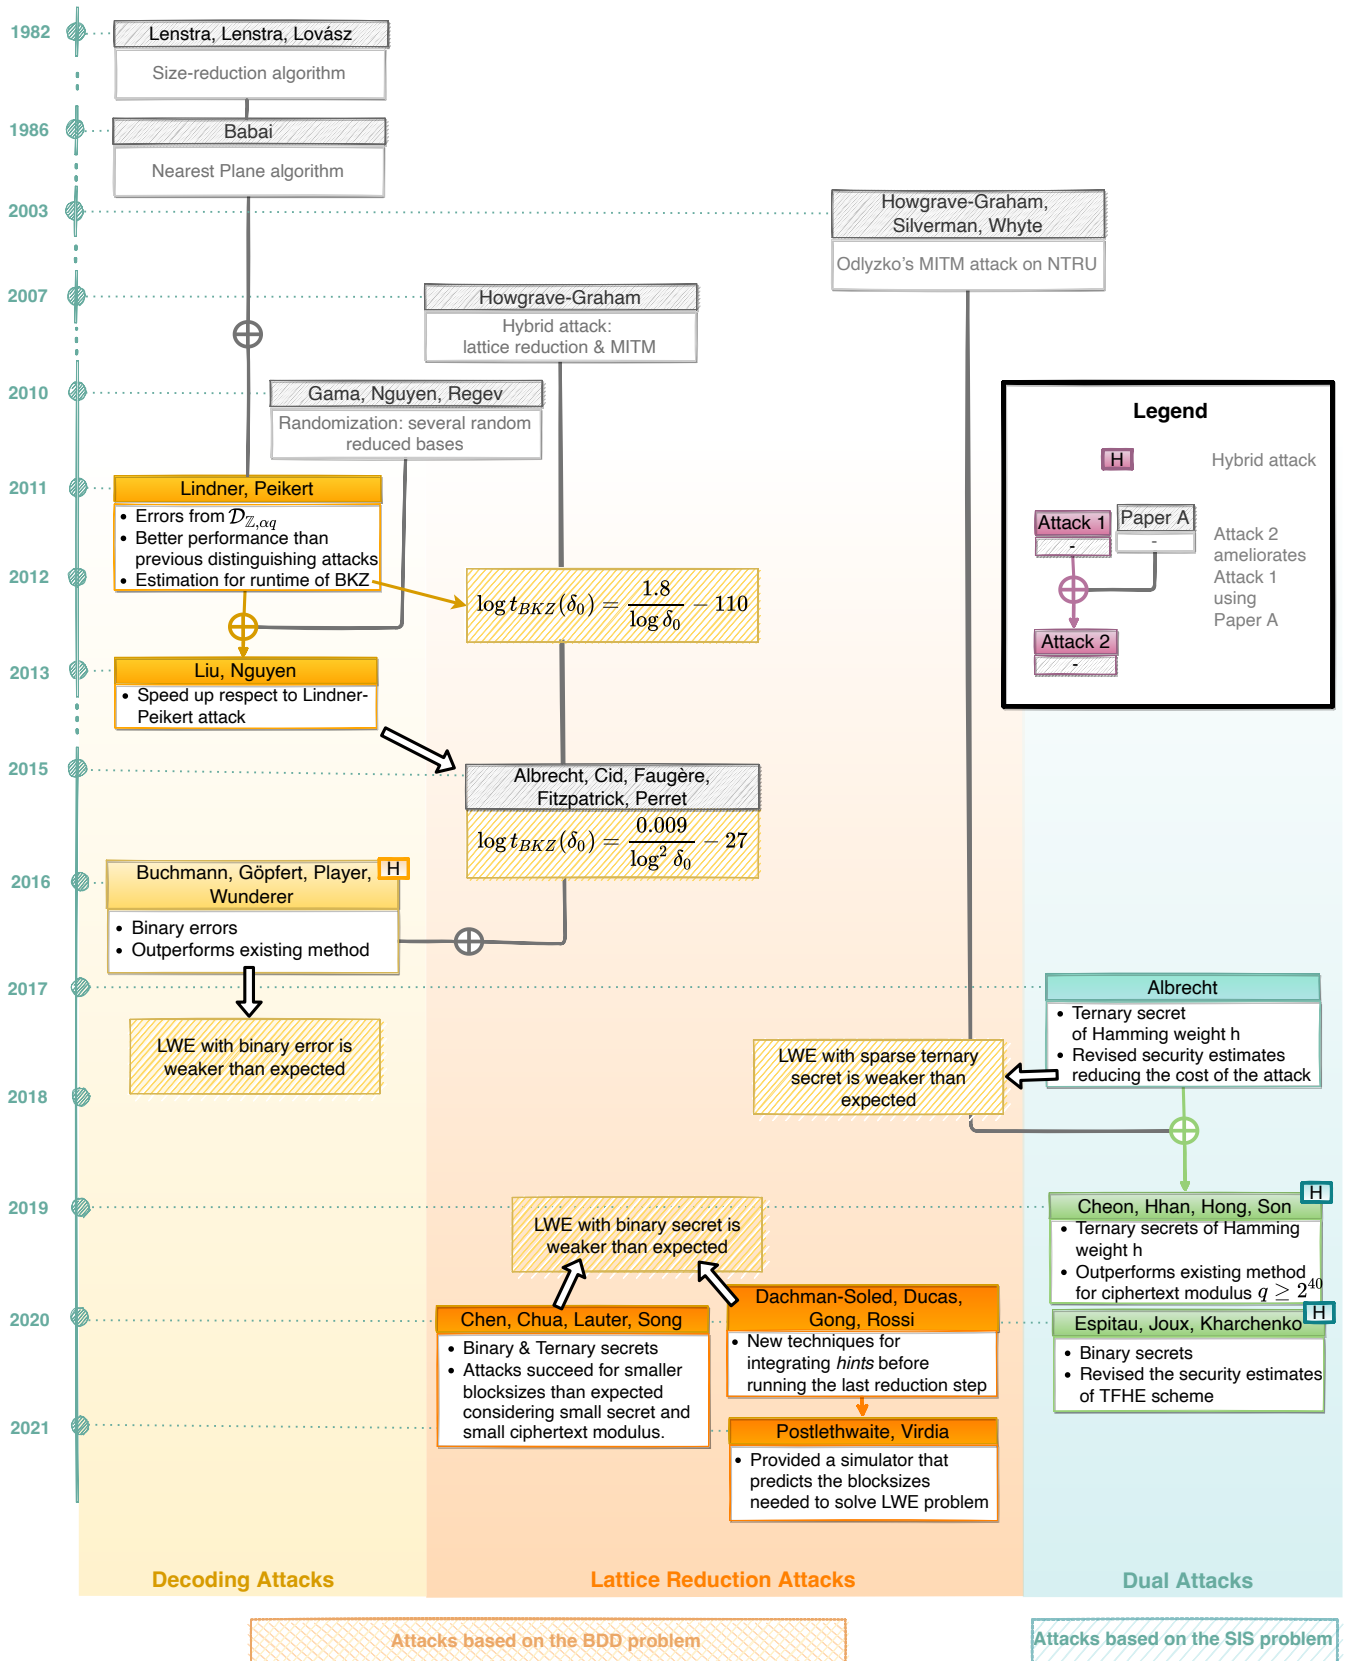


Fig. 12. Timeline of the main attacks on the LWE problem.

Subsequently, Buchmann, Göpfert, Player and Wunderer [154] provided a hybrid attack for the LWE setting, following the approach of Howgrave-Graham [155], who combined the lattice reduction and the meet-in-the-middle (MITM) attacks. The authors showed that, for specific parameters and in the binary error setting (i.e. the errors are random vectors in $\{0, 1\}^m$) their attack surpasses previous attacks on LWE. We refer to Wunderer's PhD thesis for the analysis of the hybrid decoding attack [156].

2) *Reduction of BDD problem to uSVP*: The attacks based on uSVP try to solve the LWE problem by constructing an integer embedding lattice using either the Bai-Galbraith [157] or Kannan [158] technique. The main idea is to embed the lattice $\mathcal{L}(A) = \{A\mathbf{y} : \mathbf{y} \in \mathbb{Z}_q^n\}$ into a higher-dimensional lattice $\mathcal{L}(A')$, where $A' = (A | I_m - \mathbf{b})$ and

$$\mathcal{L}(A') = \{\mathbf{x} \in \mathbb{Z}_q^{n+m+1} : A'\mathbf{x} = \mathbf{0} \bmod q\}.$$

This new lattice $\mathcal{L}(A')$ has dimension $d = n + m + 1$ (note that the dimension of $\mathcal{L}(A)$ is n) and a unique shortest vector $\mathbf{v} = (\mathbf{s}, \mathbf{e}, 1)$, where \mathbf{e} is the error and \mathbf{s} is the secret of the LWE instance $(A, A\mathbf{s} + \mathbf{e})$. Thus, finding the shortest vector \mathbf{v} in $\mathcal{L}(A')$ is equivalent to solving the LWE problem. The advantage of adopting the lattice $\mathcal{L}(A')$ is that finding the shortest vector in $\mathcal{L}(A')$ is tractable, whereas in $\mathcal{L}(A)$ it is not since we do not know whether $A\mathbf{s}$ is a shortest vector in $\mathcal{L}(A)$.

There are two methods to estimate the cost for solving LWE using the uSVP strategy. The first one is proposed by Gama and Nguyen [146] (called *2008 estimate*) and later updated by Albrecht, Fitzpatrick and Göpfert [159]. The second method (called *2016 estimate*) is given in [160] by Alkim, Ducas, Pöppelmann and Schwabe, where the authors predicted that \mathbf{e} can be found if

$$\sigma\sqrt{b} \leq \delta_0^{2b-d-1} \cdot q^{m/d},$$

where σ is the standard deviation of the error distribution, b is the block size of the underlying lattice reduction algorithm and δ_0 is the root-Hermite factor (Section III-A). After that, Albrecht, Göpfert, Virdia and Wunderer [161] compared these two estimates and verified experimentally the prediction of [160] when the error vector was sampled coefficient-wise from a discrete Gaussian distribution. In 2019, Bai, Miller and Wen [162] revisited the previous analysis of [160], [161], and provided experiments on estimating the cost of solving LWE via the uSVP suggesting that the 2016 estimate has higher accuracy than the 2008 estimate.

In a recent work, Dachman-Soled, Ducas, Gong and Rossi [163] generalized the uSVP attack and proved that the predictions of [160], [161] are not accurate for small block sizes (i.e. $b \leq 30$). In a parallel work, Chen, Chua, Lauter and Song [164] showed similar results. Namely, they found that for settings with a small error and secret values, i.e. sampled from a binary and ternary distributions, the 2008 and 2016 estimations [160] are optimistic for a small block size (i.e. $b \leq 45$). These studies show that the security levels for small block sizes are smaller than initially described by the online

LWE Estimator [130]. However, it is important to mention that as Dachman-Soled, Ducas, Gong and Rossi [163] confirm, small block sizes are not relevant. In 2021, Postlethwaite and Virdia [165] improved the result of [163] and provided a simulator that predicts the block sizes needed to solve uSVP instances via lattice reduction.

C. Attacks Based on the Short Integer Solution Problem

The attacks based on the dual strategy (also called *dual attack*) consist in solving the LWE problem via the SIS strategy (Section III-F), namely, of finding a short vector in the scaled dual lattice $\mathcal{L}_q^\perp = \{\mathbf{x} \in \mathbb{Z}_q^m : \mathbf{x}A \equiv \mathbf{0} \bmod q\}$. Note that this problem is equivalent to solve the Decision-LWE problem. Indeed, given LWE samples (A, \mathbf{b}) , we can decide whether $\mathbf{b} = A\mathbf{s} + \mathbf{e}$ or \mathbf{b} is uniformly random by computing $\langle \mathbf{v}, \mathbf{b} \rangle$ where \mathbf{v} is the short vector in the lattice \mathcal{L}_q^\perp . In fact, if \mathbf{b} is not random, i.e. $\mathbf{b} = A\mathbf{s} + \mathbf{e}$, we have

$$\langle \mathbf{v}, \mathbf{b} \rangle = \langle \mathbf{v}A, \mathbf{s} \rangle + \langle \mathbf{v}, \mathbf{e} \rangle \equiv \langle \mathbf{v}, \mathbf{e} \rangle \bmod q.$$

Since $\langle \mathbf{v}, \mathbf{e} \rangle$ is short (i.e. \mathbf{v} and \mathbf{e} are sufficiently short), the adversary has to check if $\langle \mathbf{v}, \mathbf{b} \rangle$ is close to zero modulo q .

The advantage of distinguishing $\langle \mathbf{v}, \mathbf{e} \rangle$ from random, computed by Lindner and Peikert in [147], is close to

$$e^{-\pi(\|\mathbf{v}\|\alpha)^2},$$

where α is given by the Gaussian distribution χ , namely, αq is the width parameter of χ (Section II-B). To produce such a short \mathbf{v} we require a lattice reduction algorithm. Note that, the outcome of the lattice reduction is a vector $\|\mathbf{v}\| \approx \delta_0^m q^{n/m}$, but δ_0 depends on the algorithm used, and Micciancio and Regev [26] showed that the minimum for $f(m) = \delta_0^m q^{n/m}$ is obtained when $m = \sqrt{n \log q / \log \delta_0}$.

Albrecht in [134] presented a variant for the dual attack taking into consideration small and sparse secrets. Also, Cheon, Hhan, Hong and Son [166], proposed a new hybrid attack combining the dual attack of Albrecht [134] and the MITM attack on NTRU by Howgrave-Graham, Silverman and Whyte [167]. This hybrid attack outperforms the dual attack for some specific parameter sets of the homomorphic encryption scheme, namely, for sparse ternary secrets, but it was extended by Espitau, Joux and Kharchenko [168] to binary secrets as well.

D. Exhaustive Search on Secret Key \mathbf{s}

This strategy consists of directly finding \mathbf{s} such that $\|A\mathbf{s} - \mathbf{b}\|$ is small. This can be achieved by performing the Arora-Ge algorithm [169]. This algorithm uses a linearization technique that mainly consists of adding new variables in the system to transform non-linear into linear equations. It also adopts the assumption that the error lies in a fix range. The Arora-Ge algorithm solves the LWE in time $2^{\tilde{O}(n^{2\varepsilon})}$, where ε is such that $\alpha q = n^\varepsilon$ and αq is the width parameter of the Gaussian distribution χ (Section II-B).

E. Attacks on the RLWE Problem

As previously mentioned, RLWE-based schemes are equally secure as the LWE version when the error distribution is correctly chosen. However, there are known examples of error distributions that are insecure for certain rings. In 2015, Elias, Lauter, Ozman and Stange [170] provided an attack on the decision version of the RLWE problem for two specific families of polynomial functions (namely, the definition of the polynomial considers the ciphertext modulus of the scheme). In [171], Chen, Lauter and Stange generalized this attack to certain Galois number fields and defined a new solution for the RLWE problem. These papers were later improved by Chen, Lauter and Stange in [172] and by Castryck, Iliashenko and Vercauteren in [131], [132].

Note that in the Homomorphic Encryption Security Standard paper [55], the authors provide secure parameters for RLWE schemes over power-of-two cyclotomic rings. On the other hand, for generic cyclotomic rings, Ducas and Durmus [51], Lyubashevsky, Peikert and Regev [50], [173] and Crockett and Peikert [174] investigated the types of the error distribution and proposed different ways of choosing a safe error polynomial.

F. Concrete Parameters

Finding an optimal set of parameters for a specific FHE scheme is challenging since it is function dependent. For example, for the second generation schemes, the complexity (i.e. depth) of the function to be homomorphically evaluated impacts the error growth. Higher depths, require higher ciphertext modulus q , and the adoption of a higher modulus decreases the security level. The security level can be increased by adopting a higher polynomial degree, but this impacts efficiency. Some works [92], [93] have proposed theoretical bounds for error growth estimation, which can be used to obtain the parameters heuristically. However, these works are too conservative with respect to the parameters used in practice. The reason for this is that these theoretical bounds seek for very low failure probability (e.g., less than 2^{-55}), whereas in practical scenarios smaller values are still probabilistically acceptable. The main open problem in the field of parameter selection is that there is a significant gap between the parameters obtained theoretically using the previously proposed heuristics and the parameters used in practice and obtained in a trial an error fashion [92].

This question is addressed by the Homomorphic Encryption Security Standard [55] presents some recommended (and conservative) parameters for FHE schemes, following the LWE Estimator [130]. Specifically, starting from the dimension $n = 2^k$ (with $k = \{10, \dots, 15\}$), the authors of [55] cater for recommended values of the q , for a given security level $\lambda \in \{128, 192, 256\}$. In this standard, the error follows a discrete Gaussian distribution with standard deviation $\sigma \approx 3.2$, whereas the distribution for the secret key can be:

- uniform ternary, i.e. the secret s is chosen uniformly at random from $\{-1, 0, 1\}^n$;
- uniform, i.e. the secret s is chosen uniformly at random from \mathbb{Z}_q^n ;

- Gaussian with $\sigma \approx 3.2$, the same as the error distribution.

It is worth commenting that, despite the valuable contribution and the impact of the Homomorphic Encryption Security Standard [55], there are some limitations, which have been pointed out by Curtis and Player [175], namely:

- 1) The standard does not consider a sparse ternary secret of Hamming weight h (i.e. a distribution where the elements are sampled uniformly at random from $\{-1, 0, 1\}^n$ with exactly h components different from zero). Note that many implementations use exactly this secret distribution (e.g. HEAAN uses sparse ternary secret with $h = 64$).
- 2) The consideration of sparse secrets is not included in the standard since there exists a wider range of attacks that can be applied [175].
- 3) The standard, as well as the LWE Estimator, does not consider hybrid attacks. In particular, when the secret vector follows a sparse distribution with Hamming weight $h = 128$, hybrid attacks are very powerful and, as a result, we have a noticeable security loss [175]. In fact, the new version of the Estimator, i.e. the Lattice Estimator, considers this kind of attacks.

To conclude this section we highlight the weakness of adopting secrets from a sparse distribution and the deleterious effects of hybrid attacks with one example reported by [175, Table 2] in Table II. Specifically, in Table II, λ_{target} is the currently standardised LWE security level for specific $n = 2^{10}$ and q for a uniform ternary secret [84]. The last four columns represent the security of each parameter set against uSVP attacks (Section VI-B2), dual attacks (Section VI-C), hybrid decoding attacks [154] (Section VI-B1), and hybrid dual attacks [166] (Section VI-C) for a *sparse* secret with Hamming weight $h = 128$, using the BKZ algorithm. Note that [175] used a conservative analysis for both hybrid attacks.

TABLE II
SECURITY LEVEL IN BITS FOR A SPARSE TERNARY SECRET OF HAMMING WEIGHT $h = 128$ AND HYBRID ATTACKS AS IN [175, TABLE 2].

n	λ_{target}	$\log q$	uSVP	dual	hybrid-dec	hybrid-dual
1024	128	27	124.9	127.8	111.5	106.2
	192	19	178.2	178.8	146.2	141.8
	256	14	235.5	238.5	181.5	176.6

VII. FHE FOR MACHINE LEARNING

This section provides a comprehensive view on the combined topic of privacy-preserving and machine learning, which, though fairly new, has been the subject of multiple research efforts. Machine Learning (ML) refers to a set of algorithms and computing systems used to build models that incorporate or learn structural knowledge of input datasets. A limitation to a wide adoption is the fact that machine learning mandates access to a large amount of data to achieve high accuracy rates, thus introducing data privacy and security concerns. FHE facilitates arithmetic evaluations of encrypted

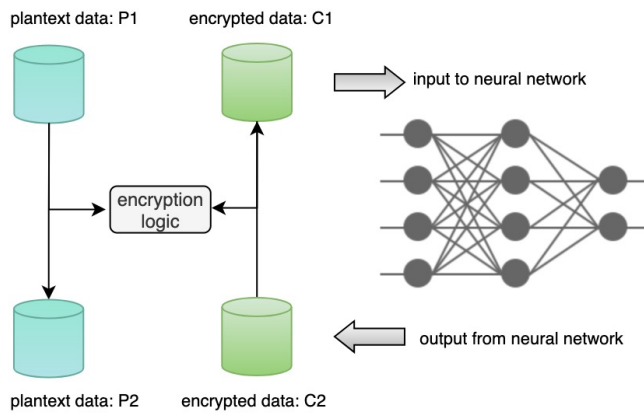


Fig. 13. Example of a potential privacy-preserving ML model.

data of real numbers, which in turn enables the development of privacy-preserving machine learning training algorithms, and potentially provides a way to overcome aforementioned privacy and security concerns. An example of a potential privacy-preserving machine learning model is shown in Figure 13.

A. Support Vector Machines

Support Vector Machines (SVMs) are widely used for their performance in classifications tasks, and multiple privacy-preserving SVM computing schemes have been proposed. [176] proposes a privacy-preserving scheme for both SVM training and classification using additive homomorphic encryption and secret sharing or secure multiparty computation protocols. [177] presents an algorithm based on Homomorphic Encryption for the SVM training phase which avoids inefficient operations within an encrypted domain. [178] proposes a two-class and multi-class classification protocol which uses SVMs, which exploits Paillier's cryptosystem [8] and secure two-party computation (client and server parties hold a share of the secret). In a more practical implementation, [179] presents EPIC, an image classification system trained with an SVM computing scheme, while input features are extracted based on the techniques of transfer learning. EPIC used Multi-Party Computation (MPC) tools to achieve privacy-preserving classification tasks and can be applied to homomorphic encryption domain.

B. Neural Networks and Other Machine Learning Models

Neural networks can be thought of as a generalization of regression to present elaborate relationships between high dimensional input data and output data. Privacy-preserving machine learning with neural networks has been addressed by multiple research efforts, even though computational complexity remains a challenge especially when neural networks are used for training over encrypted data.

Authors in [180] use training algorithms, which can be expressed as low degree polynomials, in order to train over encrypted data leveraging SHE. While this works well on very limited applications, the accuracy of the proposed system is relatively low and cannot compete with neural networks. It also

cannot be scaled to more complex operations such as division or exponentiation. [181] proposes a scheme that uses three homomorphic systems (i.e. Paillier cryptosystem, Quadratic Residuosity, and BGV scheme) and garbled circuits to provide privacy-preserving classification for three different machine learning algorithms: Hyperplane Decision, Naive Bayes, and Decision trees, where features' description is assumed public. [182] presents methodologies to train machine learning models such as random forests - using a stochastic fraction estimator - and naïve Bayes - using a semi-parametric model for class decision boundary - and demonstrate their accuracy when applied to data encrypted with homomorphic encryption. [183] presents a hardware architecture that implements Bayesian filters and Decision Trees (DT) for homomorphically encrypted data. [184] investigates multi-key FHE using collaborative learning over input datasets encrypted with different encryption schemes and keys. The approach, however, suffers from scalability issues and high computational complexity. [185] presents CryptoNets which applies a neural network - an artificial feed-forward neural network, known to a specific party and trained on plaintext data - to make predictions with a high accuracy on homomorphically encrypted data. The performance of CryptoNets is rather limited due to the replacement of the sigmoidal activation function and the computational overhead. [186] proposes a privacy-preserving deep learning model - a double-projection deep computation model whereas learning is outsourced to a cloud layer to improve the learning efficiency - trained with a back-propagation algorithm and uses BGV scheme. Improving on CryptoNets, [187] presents a version of this latter which improves latency and memory usage. [188] presents a methodology to train a convolutional neural network (CNN) model using homomorphically encrypted data, yielding high performance overheads. [189] presents a CNN used for image classification with FHE properties on Graphics Processing Units (GPU), to accelerate classification, while maintaining a high accuracy rate.

C. The Industry Role

In addition to research efforts, multiple commercial products are being proposed to solve real-world problems across industry verticals. Zama's open source technology [190] enables trained machine learning models, regardless of the underlying architecture or training method, to run inference on encrypted user data using homomorphic encryption. Application of this technology could be extended to the medical field, image classification, autonomous environments and smart cities data processing. Intel [191] and Ant Group [192] have announced a joint effort [193] to build Privacy-Preserving Machine Learning (PPML) on top of Intel's Software Guard Extensions (SGX) and Occlum, Ant Group's memory-safe, multi-process library operating system for Intel SGX, using cryptographic technologies such as homomorphic encryption and differential privacy. Duality Technologies [194] is a company providing privacy-preserving data collaboration platforms using homomorphic encryption. It has been chosen by DARPA along with other top research institutes to accelerate the use of FHE as part of DARPA's Data Protection in Virtual Environments

(DPRIVE) program, which seeks to develop a hardware accelerator for FHE computations [195].

D. Research Directions

While considerable advances have been achieved, privacy-preserving neural networks using homomorphic encryption still suffer from high computational complexity, low efficiencies, and inadequacy of deployment in real world scenarios. Further research is required to develop efficient frameworks enabling training and evaluation of complex neural networks over encrypted data or encrypted neural networks trained over plaintext data. Research directions could include:

- Algorithmic improvements: this includes (i) usage of pre-trained models to reduce computational complexity during the training phase, (ii) approximation of activation functions using polynomials, etc.
- Hardware acceleration: this includes parallelization and partitioning of the implementation of privacy-preserving models using homomorphic encryption and inherent operations on GPU cores (including hybrid CPU-GPU architectures), FPGAs, ASICs, and reconfigurable processors.

VIII. HE IN FOG COMPUTING FOR IOT

Fog computing was initially proposed by CISCO to support scalable massive IoT deployments [196] [197], but a similar concept has been adopted in 5G/6G cellular networks and referred to as edge computing [198]. It defines a layer between the IoT device and the cloud service, as close as possible to the device, where data is pre-processed. Pushing pre-processing operations close to the device is paramount to reduce both bandwidth consumption and the latency of IoT applications. In this scenario, HE can provide the missing privacy feature since pre-processing tasks can be done over encrypted data. However, fog computing also presents intrinsic characteristics that must be taken into account at the time of applying HE. Unlike cloud computing, fog computing considers data in motion, i.e. moving through the network at the generation rate of each specific device. Data processing is event-driven (triggered by the device) and performed packet-by-packet. Thus, data processing is delay intolerant and the scope of the processing tasks is limited to the information contained in a single data packet: relevance/category evaluation; formatting; encoding; expanding/compressing; filtering; or assessing thresholds and real-time alerts.

The Smart City scenario provides an example of the applicability of fog computing and how HE can become relevant. Smart cities cover a wide set of applications [199] such as intelligent transportation, efficient resource distribution (lighting, water and waste management), safety and security, or environmental monitoring. These applications have a common requirement, they are supported by massive IoT deployments composed of small sensors constantly fuelling data to smart city data collectors. The fog layer is in charge of pre-processing the data in intermediate gateways, which is fundamental to maintain scalability [200]. HE can be adopted to provide privacy preservation for citizens. Specifically, sensors can encrypt the data with the data collector's public key, and

the fog layer processing operations can be performed over encrypted data. Only the data collector can decrypt the data with its secret key. This approach not only prevents data disclosure in fog nodes, but also provides privacy for IoT nodes with respect to the data collector since data samples are aggregated [201], [202].

Despite its potential, HE presents a difficult fit in IoT due to its computational complexity and the ciphertext expansion. The latter refers to the fact that a ciphertext is far larger than the corresponding plaintext, hence it adds a considerable communication overhead. The former incurs a computational delay at the time of acquiring and transmitting data samples. This is aggravated by the limited hardware capabilities and computational constraints of IoT devices, and the bandwidth constraints of current IoT communication standards. Aforementioned constraints have fostered research on hybrid protocols combining HE and symmetric key encryption (SKE). In hybrid homomorphic encryption (HHE), the IoT device encrypts data using a SKE scheme, with a randomly generated key, and then encrypts this key with a HE scheme using the data collector's public key. A SKE scheme is less complex and is not affected by ciphertext expansion. The intermediate fog nodes can homomorphically evaluate the decryption circuit of the SKE scheme and convert SK-encrypted data into HE-encrypted data. Then, the data can be processed homomorphically and sent to the data collector, see Fig. 14. The advantage is that the ciphertext expansion and the complexity is moved from the IoT devices to the fog nodes. The IoT device is only required to encrypt homomorphically a short key. Several works have proposed AES as a SKE scheme [77] [73], since hardware acceleration for AES is a common feature in modern chips integrated in IoT devices. However, AES decryption circuit requires a multiplicative depth of at least 40, which increases the complexity in fog nodes. Some other approaches suggest public key encryption (PKE) instead of SKE [203], since the multiplicative depth is lower. However PKE has higher ciphertext expansion and complexity than SKE. In this framework, low-depth symmetric key ciphers such as [204], [205], [206] can provide noticeable gains.

Fog computing is not limited to the Smart City scenario. New emerging concepts such as Industry 4.0 or eHealth, which have coined the terms of Industrial IoT (IIoT) and Internet of Medical Things (IoMT), will also depend on the feasibility of deploying scalable micro-sensor systems. These scenarios will impose even more strict privacy requirements that could be potentially solved with HE. In general, hybrid HE should be envisioned as a solution for privacy-preserving data aggregation. However, it is worth commenting that hybrid protocols are effective when data is encrypted with the receiver's public key (the data collector in the smart city scenario). In a scenario where the IoT device encrypts with its own public key and the communication is bidirectional (the IoT device receives the encrypted processed data), the ciphertext expansion problem is unavoidable.

IX. HE IN CLOUD COMPUTING

Homomorphic encryption can become a cornerstone component for technologies within the 5G/6G realm, namely for

(according to its standard definition), some can be CCA1-secure [216]. This is not true for HPRE, all known CCA1-secure HPRE schemes are only partially homomorphic.

B. Homomorphic Authenticated Encryption

In some scenarios, privacy is not sufficient. The user may pay for a specific service [217], or use remote data processing for safety-critical applications [218]. Thus, a guarantee that the data has been processed correctly by the cloud service may be required. Note that, even if the cloud service is not malicious, it maybe be willing to submit wrong data to avoid the heavy computational load of processing homomorphically encrypted data. In this scenario, the user should be able to verify that the decrypted data is the result of a specific arithmetic circuit over the transmitted encrypted data. Fortunately, this feature can be achieved with homomorphic authenticated encryption (HAE). HAE can be obtained by composing HE and homomorphic authentication (HA) [219], [220]. Specifically, the user sends the ciphertexts and attaches homomorphic authenticators in the form of homomorphic signatures (HS). These signatures can be evaluated homomorphically, similarly to ciphertexts, to produce a valid signature for the processed data (see Fig. 16). In fact, composing HE and HA caters for the interesting property that if both the HE and HA schemes are CPA secure then the resulting HAE scheme is CCA1 secure [219].

HS were initially proposed for linear arithmetic circuits [221], [222], [223], [224]. Scenarios such as secure random linear network coding (RLNC) adopted HS [225] to counteract tag pollution attacks, although a symmetric-key based solution such as homomorphic MACs [226] can also fit in RLNC. Subsequent works provided HS schemes that accept polynomial homomorphic operations [227], [228], [229]. The works in [230], [231] provided the desirable feature of being fully homomorphic (although the former is only leveled fully homomorphic), hence enabling fully homomorphic authenticated encryption (FHAE). However, previous HS schemes (except for [230]) are selectively secure (i.e. the attacker is only provided with signatures of chosen messages before the challenge is available). Adaptive security was first achieved in [232] and by subsequent works such as [230], [233]. Unfortunately, previous HS constructions have a limitation in terms of efficiency for circuits of polynomial depth [232]. Another solution, potentially more efficient, is the adoption of verifiable computation (VC) schemes that work over encrypted data [234], [235], [236]. A VC scheme provides a proof that each arithmetic gate of the arithmetic circuit has had its inputs processed. Moreover, it is even possible to provide such a proof of computation over a partially private circuit (known only to the cloud service), since the cloud service could prove that part in zero-knowledge.

C. Homomorphic Encryption in Multi-Party Computation

HE provides a solution for the centralization of private computations. However, in a scenario where several parties aim to interact, the direct application of HE is not so intuitive. Specifically, several cloud services may want to evaluate a function combining their private datasets without leaking any

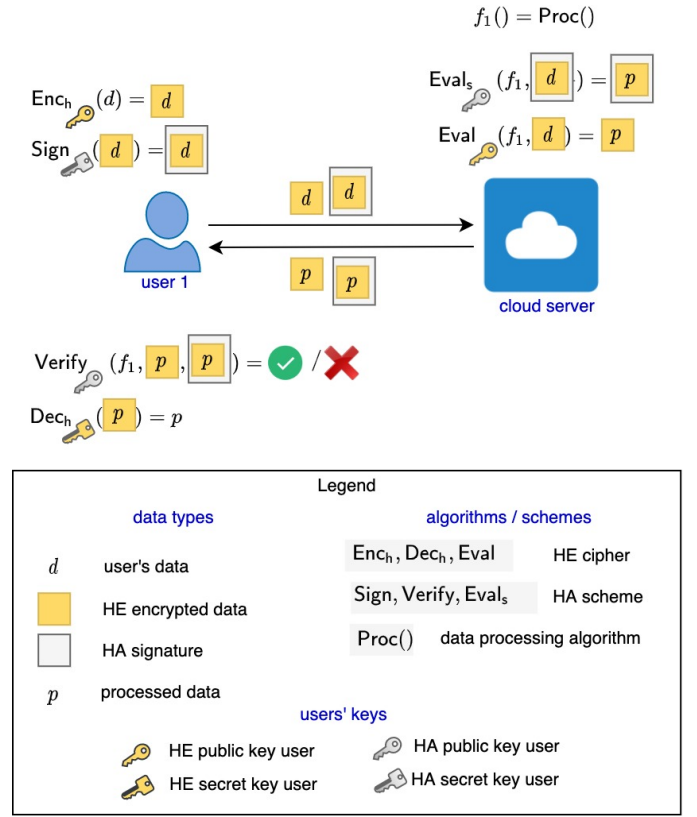


Fig. 16. FHAE by composition encrypt-then-sign.

information about the inputs (except for what can be inferred from the output). Such scenario can be addressed with secure multi-party computation (MPC) [237] [238] [239]. There are different kinds of MPC protocols optimized for arithmetic and boolean circuits, based on secret sharing techniques [240], [241] and garbled circuits [242], [243] respectively. Interestingly, some of these protocols follow a pre-processing model where the computation is divided into two phases. The first phase happens before the parties' inputs are defined, and consists of the generation of cryptographic material (secret-shared elements or garbled circuits) that is later consumed to speed up the second phase. In the second phase, parties define their inputs and evaluate the circuit privately. The concept of consuming elements refers to the fact that this material cannot be used twice, hence it must be generated for each execution. It is precisely in the generation of the pre-processing material where HE still plays a fundamental role in MPC. Protocols like [243], [241], [244] use leveled HE. It is worth commenting that, in this context, HE could be replaced by oblivious transfer (OT). In some protocols OT is more efficient than HE, but this must be evaluated per individual cases. Specifically, the protocol in [241] adopts a HE-based approach that improves the OT-based version of the same protocol in [245]. However, [246] adopts OT and is more efficient than [243], which uses SHE. The key piece that makes HE faster than OT in [241] is the existence of an efficient zero-knowledge proof to prove knowledge of a plaintext in a HE-encrypted ciphertext, which is required to provide active security. On the other

hand, [246] adopts OT because the OT protocol is compatible with an optimization generally adopted in garbled circuits (the FreeXOR technique).

Although less efficient than previous approaches, MPC can also be constructed directly with Multi-Key FHE (MKFHE). The work in [247] proposes a construction that requires only two communication rounds, see Fig. 17. In the first round, each party encrypts its inputs under multiple keys and broadcasts the ciphertexts to all parties. Then, each party evaluates the circuit homomorphically. Finally, in the second round, each party partially decrypts the result and broadcasts its share of the output. All shares can be combined locally by each party to obtain the output. In MKFHE, a ciphertext cannot be decrypted without all partial decryptions from the secret key holders. Hence, the input privacy is guaranteed. However, this also means that if only one party fails in delivering its share of the output, the MPC protocol would fail. This was addressed in [248] with a three-round MPC protocol that adopts Threshold MKFHE (TMFHE). TMFHE enables a smaller subset of parties to decrypt a ciphertext. Hence, after an initial round of input sharing, any subset of sufficient size can reconstruct the output. This makes the protocol resilient to failures, but it also requires trust since smaller subsets can decrypt the parties' private inputs. It is worth commenting that MKFHE is sufficient to provide MPC with passive security, i.e. malicious parties follow the protocol specification and try to extract information from the transmitted messages. But in a setting where parties can deviate from the protocol specification (e.g. transmit wrong shares of the output) active security is needed, and can be achieved by integrating zero-knowledge proofs.

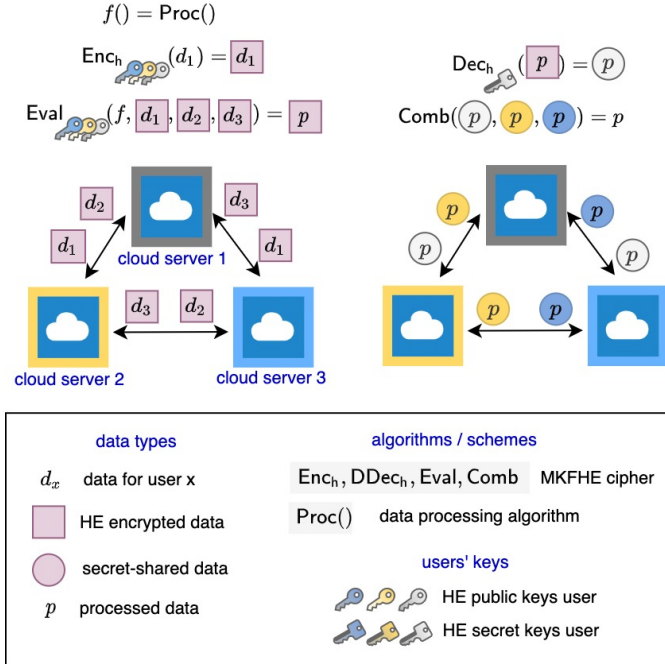


Fig. 17. MPC constructed with MKFHE. Only the operations performed by cloud server 1 have been detailed in the figure, but the rest of the parties behave analogously. For simplicity, the algorithm Enc in the figure generates a ciphertext for multiple keys, but in [247] this is achieved with an *Expand* algorithm.

X. TOWARDS PRACTICAL FHE-BASED APPLICATIONS

The first FHE schemes were about 10^9 times slower than plaintext computations [2], and were hence considered far from being practical. Optimizations achieved over the past decade have tremendously improved the performance of FHE schemes [28].

From a software perspective, FHE libraries have been pivotal in helping researchers and practitioners write the first FHE-based applications, and their evolution and optimization has significantly increased the efficiency of such applications over the past years. However, utilizing such APIs requires deep knowledge of FHE schemes. Recently, higher-level tools have evolved, attempting to bridge the gap between engineers developing privacy-preserving applications and the technical FHE libraries at hand. Along the same lines, a number of FHE compilers have become available with the objective of converting high-level programs to FHE-based implementations. These compilers are a key step towards making FHE available to non-experts that require such foundational blocks to design privacy-preserving applications, hence contributing to broad FHE adoption.

From a hardware point of view, considerable efforts leveraging specific hardware architectures (e.g., FPGA and ASIC-based) have been made. Such designs are referred to as FHE accelerators, and provide substantial improvements of FHE schemes performance in software.

The rest of this section presents an overview of the most relevant software and hardware works proposed in literature, discusses trade-offs observed in terms of performance and other metrics, and describes few real-world applications currently leveraging FHE in production systems.

A. FHE Libraries

The principal objective of FHE libraries is to make FHE scheme operations available via an API. Besides the core functionality provided by KeyGen, Enc, Dec and Eval, most of the widely used libraries incorporate additional features that allow ciphertext maintenance (i.e. noise growth management during computations) and manipulation, as well as homomorphic addition and multiplication methods. However, the correct utilization is left for the developers who must have an in-depth knowledge of what each API call entails in a given privacy-preserving solution.

Table III provides available open source FHE libraries, the language in which they are written, supported FHE schemes, and the date of the last update release. The first library ever published is HELib (Homomorphic Encryption Library), by Halevi and Shoup [76], [259], which is implemented in C++ and built on top of the NTL library [260]. SEAL (Simple Encrypted Arithmetic Library), which is developed by Microsoft [80], is implemented in C++ and C# (to support .NET). It utilizes Intel's HEXL [261], a library providing efficient implementations of homomorphic encryption operations, specifically targeting AVX512-enabled processors. PALISADE is developed by Duality Technologies and the New Jersey Institute of Technology [249]. It is written in C++ and can be configured to use the NTL library. Lattigo [250] was proposed

TABLE III
OPEN SOURCE LIBRARIES FOR FHE SCHEMES

Library	Language	Scheme					Date of last commit
		BGV	B/FV	FHEW	TFHE	HEAAN	
HElib [76]	C++	●	○	○	○	●	8/9/2021
SEAL [80]	C++/C#	○	●	○	○	●	22/9/2021
PALISADE [249]	C++	●	●	●	●	●	17/9/2021
Lattigo [250]	Go	○	●	○	○	●	19/9/2021
FHEW [251]	C++	○	○	●	○	○	30/5/2017
TFHE [252]	C++/C	○	○	○	●	○	16/9/2021
concrete [253]	Rust	○	○	○	●	○	20/9/2021
HEAAN [254]	C++	○	○	○	○	●	2/9/2021
RNS-HEAAN [255]	C++	○	○	○	○	●	26/10/2018
FV-NFLlib [256]	C++	○	●	○	○	○	26/7/2016
CuFHE [257]	Cuda/C++	○	○	○	●	○	9/2/2019
NuFHE [258]	Python	○	○	○	●	○	18/3/2020

by Mouchet, Bossuat, Troncoso-Pastoriza and Hubaux [262] and is the first library written in Go. The FHEW library [251] by Ducas and Micciancio, is written in C++, however, has not been updated since 2017. The TFHE library [252] was provided by the authors of the TFHE paper [13], is written in C++ and C, and requires at least an Fast Fourier Transform (FFT) processor to run. TFHE is considered to be the successor of the FHEW library. Concrete [253], [263] is Zama’s variant of TFHE implemented in Rust. The HEAAN library [254] is implemented in C++ and is built on top of the NTL library. The RNS-HEAAN library [255] by Kyoohyung and Miran is implemented in C++; it has not been updated since 2018. FV-NFLlib [256] is written in C++ and is based on the NFLlib C++ library [264]. NFLlib is a library dedicated to ideal lattice-based cryptography, and is based on the Number Theoretic Transform (NTT). It is important to note that FV-NFLlib has not been update in the last 5 years. Finally, the CuFHE [257] and NuFHE [258] are two GPU-based libraries that implement TFHE in CUDA. Specifically, the CuFHE library adopts an implementation of NTT, GPU-accelerated, that is based on [265] by Dai and Sunar. The NuFHE library provides support for either FFT or purely integer NTT.

B. FHE Compilers

FHE compilers are high-level tools that aim at abstracting the technical APIs exposed by FHE libraries, so that a wider range of developers are able to implement privacy-preserving mechanisms securely. As noted by Viand, Jattke and Hithnawi [266], FHE compilers tackle some of the most common engineering challenges that exist nowadays when designing FHE-based applications:

- Parameters choice: defining appropriate parameter values for FHE schemes resulting in secure and efficient instances is not a simple task. Some FHE compilers allow for some sort of automatic parameter generation according to some predefined requirements.
- Plaintext encoding: in FHE, the semantics of the plaintext message are strictly related to the type of homomorphic computations that can be conducted. Some context-specific FHE compilers can already be used to aid in this particular item (e.g. nGraph-HE).

- Data-independent execution: given that FHE operations are data-independent by nature, it is not trivial to conduct data-dependent branching steps using FHE because they can break privacy properties. In this case, it is possible to have branching operations by means of evaluating both branches and selecting the result at the end.
- Packing or batching: FHE schemes allowing for message packing or batching into a single ciphertext can directly leverage SIMD instruction sets. Some FHE compilers already actively optimize for vectorized operations.
- Ciphertext maintenance: optimally managing how noise grows during FHE operations is not straightforward, and FHE compilers are starting to use advanced strategies to assist in this traditionally complicated part.

Table IV presents a list of FHE compilers showing in which programming language they are written, what FHE libraries (from the ones previously highlighted in this work) they utilize as well as the date of their latest released update. ALCHEMY [267], [277], [278] is a compiler written in Haskell by Crockett, Peikert and Sharp. It implements BGV utilizing $\Lambda \circ \lambda$ [279] (pronounced “L O L”), a library for ring-based lattice cryptography that supports also FHE. Cingulata (previously called Armadillo) [268], [280] is a compiler written in C++ by Carпов, Dubrulle and Sirdey. It is built on top of the FLINT [281] and Sage [282] libraries. The Encrypt-Everything-Everywhere (E³) [269], [283] is a framework presented by Chielle, Mazonka, Tsoutsos and Maniatakos. It is mainly written in C++ and supports a variety of FHE libraries. The Encrypted Vector Arithmetic Language and Compiler (EVA) [270], [284] was presented by Dathathri et al., is written in C++, and incorporates CHET [274] to support tensor circuits. Marble [271], [285] is a C++ compiler written by Viand and Shafagh, and RAMPARTS [272] is a compiler written in Julia by Archer et al. The Transpiler [273], [286] is a C++ tool developed by Gorantala et al. that is currently leveraging one FHE library. The nGraph-HE [275], [287] compiler by Boemer, Lao, Cammarota and Wierzynski is based on Intel’s nGraph ML compiler [288]. Support for non-polynomial activation functions was added subsequently [289]. Finally, SEALion [276] was presented by van Elsloo, Patrini and Ivey-Law. It is important to highlight that CHET, nGraph-HE and SEALion are domain-specific FHE compilers designed

TABLE IV
PUBLICLY AVAILABLE FHE COMPILERS

Compiler	Language	Library						Date of last commit
		HElib	SEAL	PALISADE	FHEW	TFHE	HEAAN	
ALCHEMY [267]	Haskell	○	○	○	○	○	○	15/3/2020
Cingulata [268]	C++	○	○	○	○	●	○	7/12/2020
E ³ [269]	C++	●	●	○	●	●	○	26/3/2021
EVA [270]	C++	○	●	○	○	○	○	1/5/2021
Marble [271]	C++	●	●	○	○	○	○	23/12/2020
RAMPARTS [272]	Julia	○	○	●	○	○	○	-
Transpiler [273]	C++	○	○	○	○	●	○	23/8/2021
CHET [274]	C++	○	●	○	○	○	●	-
nGraph-HE [275]	C++	○	●	○	○	○	○	8/7/2021
SEALion [276]	C++	○	●	○	○	○	○	-

particularly for ML applications. To the best of the authors knowledge, we note that the libraries without a date provided for their last committed update have not been found publicly. For more details about the FHE libraries and compilers, we would like to refer the readers to [266] by Viand, Jattke and Hithnawi.

C. FHE Accelerators

Previous sections presented state-of-the-art software tools that are proving crucial in the wider adoption of FHE for developing privacy-preserving solutions. Although such tools have enabled a significant acceleration of FHE schemes, the corresponding performance still falls short with respect to plaintext computations. Therefore, FHE hardware accelerators have emerged as a practical alternative to highly optimized software implementations, thus, enabling a wider range of use cases where FHE can be utilized.

Doröz, Öztürk and Sunar [290] designed an accelerator for the Gentry and Halevy scheme [61] and were able to significantly improve run times of FHE operations. Later on, Cousins, Rohloff and Sumorok [291] designed an FPGA-based accelerator focusing on the second-generation, NTRU-based scheme, LTV. They targeted a Xilinx Virtex-7 FPGA and benchmarked the performance of the CRT (and its inverse). Their results showed an improvement of 2 orders of magnitude compared to the available reference software and CPU-based version. In [292], Roy et al. presented a RLWE-based co-processor targeting NTT optimizations and achieved considerable speeds on a Virtex-6 FPGA. Moreover, Roy et al. [293] proposed an architecture where they are able to offload operations of the FV FHE scheme to an FPGA-based accelerator. They validated their design on a Xilinx Zynq UltraScale+ FPGA and obtained an improvement of over 13 times with respect to a reference FV scheme optimized software implementation. Riazi, Laine, Pelton and Dai [294] provided a new hardware design that heavily improved the NTT operation, implemented the HEAAN scheme, and can be used with a wide range of parameter sets. They compared their proposal on two FPGA devices from Intel, namely Arria 10 and Stratix 10, with an optimized version of the SEAL library and demonstrated an improvement of more than 2 orders of magnitude. Finally, Turan, Roy and Verbaudhede [295] proposed the first accelerator to leverage FPGAs available in the Amazon AWS cloud, and achieved a 20 times improvement

with respect to the software implementation of the smart meter application that they consider in their case study.

Even though FPGA-based designs bear considerable improvements and can be run on accessible FPGA platforms, they miss an important element of FHE computations, the data movement, which, in extremely optimized designs, becomes a non-negligible bottleneck. ASIC-based designs enable the possibility to tackle potential issues caused by data movement congestion. Following this approach, Juvekar, Vaikuntanathan and Chandrakasan [296] designed Gazelle, which combines two conventional techniques, namely, homomorphic encryption and garbled circuits. A low-latency ASIC targeting a secure neural network inference running on Gazelle achieved 2-3 orders of magnitude speedups. Subsequently, Gazelle was improved (i.e. 79 times faster) by Reagen et al. with Cheetah [297]. Recently, Feldmann, Samardzic et al. have presented F1 [298], a programmable FHE accelerator that employs a wide-vector processor, which is based on a static scheduling strategy and minimizes data movement. F1 is capable of producing speedups of up to 3-4 orders of magnitude with respect to state-of-the-art software implementations, and supports BGV, HEAAN and GSW. Moreover, F1 demonstrates that ASIC-based accelerators can also be programmable, given that the same resulting hardware can accelerate a variety of programs, including multiple FHE schemes.

D. Standardization and Broad Adoption

Section X-A, Section X-B and Section X-C provide a comprehensive list of the most relevant works in software and hardware that have contributed towards making FHE more practical, hence, broadening its adoption.

However, other types of initiatives equally play a key role towards having large-scale FHE deployments. For instance, the Homomorphic Encryption open industry / government / academic consortium [55] is working on a standard for homomorphic encryption. The consortium was created in 2017 by Microsoft, IBM and Duality Technologies, and, at the time of writing this work, has more than 40 participants amongst industry, government and academia. Moreover, industrial players such as Duality Technologies [194], Zama AI [190], or Cryptolab [299] aiming at deploying FHE-based solutions, are greatly contributing to the overall ecosystem. Lastly, large tech-based companies like Intel and Google,

are starting to leverage homomorphic encryption in privacy-preserving solutions such as building PPML on top of Intel's SGX [193] and Google's Password Checkup [300], which employs Private Set Intersection.

XI. CONCLUSIONS

Homomorphic encryption has been a prolific research field over the past decade. Since Gentry's first proposed scheme in 2009, several generations of schemes have emerged, fostered by the evolution of privacy-preserving technologies. Moreover, synergies with other research fields (such as Machine Learning) and with other cryptographic protocols (such as secure multi-party computation) have increased its relevance.

Nevertheless, despite the tremendous potential of the field, current FHE schemes still present limitations that hinder their applicability within real environments. The computational complexity and ciphertexts expansion render FHE unsuitable to delay-intolerant or bandwidth-limited applications. These latter are the main impediments for FHE's widespread adoption in new generation networks.

Additionally, there are no known common schemes that encompass features offered by second, third and forth generation schemes simultaneously, which would otherwise be convenient in some scenarios, such as privacy-preserving Machine Learning. More specifically, second and forth generation schemes are equipped with packing techniques, which make them efficient for matrix multiplication, while third generation schemes are the only ones to enable efficient evaluation of non-linear functions. Moreover, second and forth generation schemes are not equipped with fast bootstrapping techniques; this limits their application to their leveled version, not their fully homomorphic version. Another limitation is the absence of thorough efforts related to noise analysis, mainly for second generation schemes. There is still a gap between theoretical bounds and the real noise growth, which increases the complexity of parameter selection. These limitations have triggered numerous research proposals, not only on new schemes and analytical studies on parameters, but also on hardware accelerators. New technologies such as memory-based computation have also been proposed for memory hungry applications, such as private deep neural networks. It is precisely the advancement on hardware acceleration that can make FHE a reality, by reducing its time complexity by several orders of magnitude: cloud computing will require more efficient implementations that cannot be achieved with software-based optimizations. The potential integration of FHE in massive IoT deployments will also depend on the ongoing research efforts on hardware. These new scenarios will define strict specifications in terms of latency, bandwidth and energy efficiency, and the FHE layers would have to meet these requirements.

ACKNOWLEDGEMENTS

This work has been partly funded by the European Commission through the H2020 projects Hexa-X (Grant Agreement no. 101015956). This work has also been partially funded by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany's Excellence

Strategy – EXC2050/1 – Project ID 390696704 – Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of Technische Universität Dresden. Marc Manzano is a member of the Intelligent Systems for Industrial Systems research group of Mondragon Unibertsitatea (IT1676-22), supported by the Department of Education, Universities and Research of the Basque Country.

REFERENCES

- [1] R. L. Rivest, L. Adleman, M. L. Dertouzos *et al.*, “On data banks and privacy homomorphisms,” *Foundations of secure computation*, vol. 4, no. 11, pp. 169–180, 1978.
- [2] C. Gentry, *A fully homomorphic encryption scheme*. Stanford university Stanford, 2009, vol. 20, no. 9.
- [3] R. L. Rivest, A. Shamir, and L. Adleman, “A method for obtaining digital signatures and public-key cryptosystems,” *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.
- [4] S. Goldwasser and S. Micali, “Probabilistic encryption & how to play mental poker keeping secret all partial information,” in *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing*, ser. STOC '82. New York, NY, USA: Association for Computing Machinery, 1982, p. 365–377.
- [5] T. ElGamal, “A public key cryptosystem and a signature scheme based on discrete logarithms,” *IEEE transactions on information theory*, vol. 31, no. 4, pp. 469–472, 1985.
- [6] J. Benaloh, “Dense probabilistic encryption,” in *Proceedings of the workshop on selected areas of cryptography*, 1994, pp. 120–128.
- [7] D. Naccache and J. Stern, “A new public key cryptosystem based on higher residues,” in *Proceedings of the 5th ACM conference on Computer and communications security*, 1998, pp. 59–66.
- [8] P. Paillier, “Public-key cryptosystems based on composite degree residuosity classes,” in *International conference on the theory and applications of cryptographic techniques*. Springer, 1999, pp. 223–238.
- [9] D. Boneh, E.-J. Goh, and K. Nissim, “Evaluating 2-DNF formulas on ciphertexts,” in *Theory of Cryptography Conference*. Springer, 2005, pp. 325–341.
- [10] C. Gentry, “Computing arbitrary functions of encrypted data,” *Communications of the ACM*, vol. 53, no. 3, pp. 97–105, 2010.
- [11] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, “(Leveled) fully homomorphic encryption without bootstrapping,” *ACM Transactions on Computation Theory (TOCT)*, vol. 6, no. 3, pp. 1–36, 2014.
- [12] J. Fan and F. Vercauteren, “Somewhat practical fully homomorphic encryption,” *IACR Cryptology ePrint Archive*, 2012.
- [13] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, “TFHE: fast fully homomorphic encryption over the torus,” *Journal of Cryptology*, pp. 1–58, 2019.
- [14] J. H. Cheon, A. Kim, M. Kim, and Y. Song, “Homomorphic encryption for arithmetic of approximate numbers,” in *Advances in Cryptology – ASIACRYPT 2017*, T. Takagi and T. Peyrin, Eds. Cham: Springer International Publishing, 2017, pp. 409–437.
- [15] Z. Brakerski and V. Vaikuntanathan, “Efficient Fully Homomorphic Encryption from (Standard) LWE,” in *IEEE 52nd Annual Symposium on Foundations of Computer Science – FOCS*, 2011, pp. 97–106, full version in <https://eprint.iacr.org/2011/344>.
- [16] —, “Fully homomorphic encryption from ring-lwe and security for key dependent messages,” in *Advances in Cryptology – CRYPTO 2011*, P. Rogaway, Ed. Berlin, Heidelberg: Springer, 2011, pp. 505–524.
- [17] M. van Dijk, C. Gentry, S. Halevi, and V. Vaikuntanathan, “Fully homomorphic encryption over the integers,” in *Advances in Cryptology – EUROCRYPT 2010*, H. Gilbert, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 24–43.
- [18] N. Aggarwal, C. Gupta, and I. Sharma, “Fully homomorphic symmetric scheme without bootstrapping,” in *Proceedings of 2014 International Conference on Cloud Computing and Internet of Things*. IEEE, 2014, pp. 14–17.
- [19] C. Gupta and I. Sharma, “A fully homomorphic encryption scheme with symmetric keys with application to private data processing in clouds,” in *2013 Fourth International Conference on the Network of the Future (NoF)*. IEEE, 2013, pp. 1–4.
- [20] T. S. Fun and A. Samsudin, “A survey of homomorphic encryption for outsourced big data computation,” *KSII Transactions on Internet and Information Systems (TIIS)*, vol. 10, no. 8, pp. 3826–3851, 2016.

- [21] R. Rothblum, "Homomorphic encryption: From private-key to public-key," in *Theory of cryptography conference*. Springer, 2011, pp. 219–234.
- [22] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti, "A survey on homomorphic encryption schemes: Theory and implementation," *ACM Computing Surveys (CSUR)*, vol. 51, no. 4, pp. 1–35, 2018.
- [23] P. Martins, L. Sousa, and A. Mariano, "A survey on fully homomorphic encryption: An engineering perspective," *ACM Computing Surveys (CSUR)*, vol. 50, no. 6, pp. 1–33, 2017.
- [24] A. Aloufi, P. Hu, Y. Song, and K. Lauter, "Computing blindfolded on data homomorphically encrypted under multiple keys: An extended survey," *arXiv preprint arXiv:2007.09270*, 2020.
- [25] D. Micciancio and O. Regev, "Worst-case to average-case reductions based on Gaussian measures," *SIAM Journal on Computing*, vol. 37, no. 1, pp. 267–302, 2007.
- [26] —, *Lattice-based Cryptography*, D. J. Bernstein, J. Buchmann, and E. Dahmen, Eds. Berlin, Heidelberg: Springer, 2009.
- [27] C. Peikert, "Lattices in Cryptography," 2013, <https://web.eecs.umich.edu/~cpeikert/lic13/>.
- [28] —, "A decade of lattice cryptography," *Foundations and Trends® in Theoretical Computer Science*, vol. 10, no. 4, pp. 283–424, 2016.
- [29] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa, "Efficient public key encryption based on ideal lattices," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2009, pp. 617–635.
- [30] SVP challenge, <https://www.latticechallenge.org/svp-challenge/>.
- [31] S. R. Kumar and D. Sivakumar, "On the unique shortest lattice vector problem," *Theoretical computer science*, vol. 255, no. 1-2, pp. 641–648, 2001.
- [32] Y.-K. Liu, V. Lyubashevsky, and D. Micciancio, "On bounded distance decoding for general lattices," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, 2006, pp. 450–461.
- [33] V. Lyubashevsky and D. Micciancio, "On bounded distance decoding, unique shortest vectors, and the minimum distance problem," in *Advances in Cryptology - CRYPTO 2009*, S. Halevi, Ed. Berlin, Heidelberg: Springer, 2009, pp. 577–594.
- [34] S. Bai, D. Stehlé, and W. Wen, "Improved reduction from the bounded distance decoding problem to the unique shortest vector problem in lattices," in *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- [35] S. Khot, "Hardness of approximating the shortest vector problem in high ℓ_p norms," *Journal of Computer and System Sciences*, vol. 72, no. 2, pp. 206–219, 2006.
- [36] W. Wen, "Contributions to the hardness foundations of lattice-based cryptography. (contributions aux fondements de complexité de la cryptographie sur réseaux)," Ph.D. dissertation, University of Lyon, France, 2018. [Online]. Available: <https://tel.archives-ouvertes.fr/tel-01949339>
- [37] M. Ajtai, "Generating hard instances of lattice problems (extended abstract)," in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, ser. STOC '96. New York, NY, USA: Association for Computing Machinery, 1996, p. 99–108.
- [38] D. Micciancio, "Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions," in *IEEE 43rd Annual Symposium on Foundations of Computer Science - FOCS*, 2002, pp. 356–365.
- [39] —, "Generalized compact knapsacks, cyclic lattices, and efficient one-way functions," *Comput. Complex.*, vol. 16, no. 4, pp. 365–411, 2007. [Online]. Available: <https://doi.org/10.1007/s00037-007-0234-9>
- [40] P. Bert, P.-A. Fouque, A. Roux-Langlois, and M. Sabt, "Practical implementation of ring-sis/lwe based signature and ibe," in *International Conference on Post-Quantum Cryptography*. Springer, 2018, pp. 271–291.
- [41] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, 2005, pp. 84–93.
- [42] —, "On lattices, learning with errors, random linear codes, and cryptography," *Journal of the ACM (JACM)*, vol. 56, no. 6, pp. 1–40, 2009.
- [43] A. Blum, M. Furst, M. Kearns, and R. J. Lipton, "Cryptographic primitives based on hard learning problems," in *Advances in Cryptology - CRYPTO '93*, D. R. Stinson, Ed. Berlin, Heidelberg: Springer, 1994, pp. 278–291.
- [44] C. Peikert, "Public-key cryptosystems from the worst-case shortest vector problem," in *Proceedings of the 41st annual ACM symposium on Theory of computing*, 2009, pp. 333–342.
- [45] D. Micciancio and P. Mol, "Pseudorandom Knapsacks and the Sample Complexity of LWE Search-to-Decision Reductions," in *Advances in Cryptology - CRYPTO 2011*, P. Rogaway, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 465–484.
- [46] D. Micciancio and C. Peikert, "Trapdoors for lattices: Simpler, tighter, faster, smaller," in *Advances in Cryptology - EUROCRYPT 2012*, D. Pointcheval and T. Johansson, Eds. Berlin, Heidelberg: Springer, 2012, pp. 700–718.
- [47] Z. Brakerski, "Fully Homomorphic Encryption without Modulus Switching from Classical GapSVP," in *Advances in Cryptology - CRYPTO 2012*, R. Safavi-Naini and R. Canetti, Eds. Berlin, Heidelberg: Springer, 2012, pp. 868–886.
- [48] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé, "Classical hardness of learning with errors," in *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, 2013, pp. 575–584.
- [49] B. Applebaum, D. Cash, C. Peikert, and A. Sahai, "Fast cryptographic primitives and circular-secure encryption based on hard learning problems," in *Advances in Cryptology - CRYPTO 2009*, S. Halevi, Ed. Springer, 2009, pp. 595–618. [Online]. Available: https://doi.org/10.1007/978-3-642-03356-8_35
- [50] V. Lyubashevsky, C. Peikert, and O. Regev, "On ideal lattices and learning with errors over rings," in *Advances in Cryptology - EUROCRYPT 2010*, H. Gilbert, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 1–23.
- [51] L. Ducas and A. Durmus, "Ring-LWE in Polynomial Rings," in *Public Key Cryptography - PKC 2012*, M. Fischlin, J. Buchmann, and M. Manulis, Eds. Berlin, Heidelberg: Springer, 2012, pp. 34–51.
- [52] A. Langlois and D. Stehlé, "Worst-case to average-case reductions for module lattices," *Designs, Codes and Cryptography*, vol. 75, no. 3, pp. 565–599, 2015.
- [53] C. Peikert, O. Regev, and N. Stephens-Davidowitz, "Pseudorandomness of ring-lwe for any ring and modulus," in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, 2017, pp. 461–473.
- [54] C. Gentry, A. Sahai, and B. Waters, "Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based," in *Advances in Cryptology - CRYPTO 2013*, R. Canetti and J. A. Garay, Eds. Berlin, Heidelberg: Springer, 2013, pp. 75–92.
- [55] M. R. Albrecht, M. Chase, H. Chen, J. Ding, S. Goldwasser, S. Gorbunov, S. Halevi, J. Hoffstein, K. Laine, K. Lauter, S. Lokam, D. Micciancio, D. Moody, T. Morrison, A. Sahai, and V. Vaikuntanathan, "Homomorphic encryption security standard," HomomorphicEncryption.org, Toronto, Canada, Tech. Rep., November 2018.
- [56] C. Gentry, S. Halevi, and V. Vaikuntanathan, "i-Hop Homomorphic Encryption and Rerandomizable Yao Circuits," in *Advances in Cryptology - CRYPTO 2010*, T. Rabin, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 155–172.
- [57] Z. Brakerski, "Fundamentals of Fully Homomorphic Encryption-A Survey," in *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 25, 2018, p. 125.
- [58] F. Armknecht, C. Boyd, C. Carr, K. Gjøsteen, A. Jäschke, C. A. Reuter, and M. Strand, "A guide to fully homomorphic encryption," *IACR Cryptol. ePrint Arch.*, vol. 2015, p. 1192, 2015.
- [59] A. Silverberg, "Fully homomorphic encryption for mathematicians," *Women in Numbers 2: Research Directions in Number Theory*, vol. 606, p. 111, 2013.
- [60] N. P. Smart and F. Vercauteren, "Fully homomorphic encryption with relatively small key and ciphertext sizes," in *International Workshop on Public Key Cryptography*. Springer, 2010, pp. 420–443.
- [61] C. Gentry and S. Halevi, "Implementing Gentry's Fully-Homomorphic Encryption Scheme," in *Advances in Cryptology - EUROCRYPT 2011*, K. G. Paterson, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 129–148.
- [62] P. Scholl and N. P. Smart, "Improved key generation for gentry's fully homomorphic encryption scheme," in *IMA International Conference on Cryptography and Coding*. Springer, 2011, pp. 10–22.
- [63] D. Stehlé and R. Steinfeld, "Faster fully homomorphic encryption," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2010, pp. 377–394.
- [64] R. Cramer, L. Ducas, C. Peikert, and O. Regev, "Recovering short generators of principal ideals in cyclotomic rings," in *Advances in*

- Cryptology – EUROCRYPT 2016*, M. Fischlin and J.-S. Coron, Eds. Berlin, Heidelberg: Springer, 2016, pp. 559–585.
- [65] J.-S. Coron, A. Mandal, D. Naccache, and M. Tibouchi, “Fully homomorphic encryption over the integers with shorter public keys,” in *Advances in Cryptology – CRYPTO 2011*, P. Rogaway, Ed. Berlin, Heidelberg: Springer, 2011, pp. 487–504.
 - [66] Y. Chen and P. Q. Nguyen, “Faster algorithms for approximate common divisors: Breaking fully-homomorphic-encryption challenges over the integers,” in *Advances in Cryptology - EUROCRYPT 2012*, D. Pointcheval and T. Johansson, Eds. Springer, 2012, pp. 502–519.
 - [67] J.-S. Coron, D. Naccache, and M. Tibouchi, “Public key compression and modulus switching for fully homomorphic encryption over the integers,” in *Advances in Cryptology – EUROCRYPT 2012*, D. Pointcheval and T. Johansson, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 446–464.
 - [68] P. W. Shor, “Algorithms for quantum computation: discrete logarithms and factoring,” in *Proceedings 35th annual symposium on foundations of computer science*. IEEE, 1994, pp. 124–134.
 - [69] J. Kim, M. S. Lee, A. Yun, and J. H. Cheon, “CRT-based Fully Homomorphic Encryption over the Integers,” *Cryptology ePrint Archive*, Report 2013/057, 2013, <https://eprint.iacr.org/2013/057>.
 - [70] J.-S. Coron, T. Lepoint, and M. Tibouchi, “Batch Fully Homomorphic Encryption over the Integers,” *Cryptology ePrint Archive*, Report 2013/036, 2013, <https://eprint.iacr.org/2013/036>.
 - [71] J. H. Cheon, J.-S. Coron, J. Kim, M. S. Lee, T. Lepoint, M. Tibouchi, and A. Yun, “Batch fully homomorphic encryption over the integers,” in *Advances in Cryptology – EUROCRYPT 2013*, T. Johansson and P. Q. Nguyen, Eds. Berlin, Heidelberg: Springer, 2013, pp. 315–335.
 - [72] K. Nuida and K. Kurosawa, “(Batch) Fully Homomorphic Encryption over Integers for Non-Binary Message Spaces,” in *Advances in Cryptology – EUROCRYPT 2015*, E. Oswald and M. Fischlin, Eds. Berlin, Heidelberg: Springer, 2015, pp. 537–555.
 - [73] J.-S. Coron, T. Lepoint, and M. Tibouchi, “Scale-invariant fully homomorphic encryption over the integers,” in *International Workshop on Public Key Cryptography*. Springer, 2014, pp. 311–328.
 - [74] J. H. Cheon and D. Stehlé, “Fully homomorphic encryption over the integers revisited,” in *Advances in Cryptology – EUROCRYPT 2015*, E. Oswald and M. Fischlin, Eds. Berlin, Heidelberg: Springer, 2015, pp. 513–536.
 - [75] Z. Brakerski and V. Vaikuntanathan, “Efficient fully homomorphic encryption from (standard) LWE,” *SIAM Journal on Computing*, vol. 43, no. 2, pp. 831–871, 2014.
 - [76] Shai Halevi and Victor Shoup, “HElib,” <https://github.com/homenc/HElib>.
 - [77] C. Gentry, S. Halevi, and N. P. Smart, “Homomorphic Evaluation of the AES Circuit,” in *Advances in Cryptology – CRYPTO 2012*, R. Safavi-Naini and R. Canetti, Eds. Berlin, Heidelberg: Springer, 2012, pp. 850–867.
 - [78] F. Maino, U. Blumenthal, and K. McCloghrie, “The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model,” RFC 3826, 2004. [Online]. Available: <https://rfc-editor.org/rfc/rfc3826.txt>
 - [79] C. Gentry, S. Halevi, and N. P. Smart, “Fully homomorphic encryption with polylog overhead,” in *Advances in Cryptology – EUROCRYPT 2012*, D. Pointcheval and T. Johansson, Eds. Berlin, Heidelberg: Springer, 2012, pp. 465–482.
 - [80] “Microsoft SEAL (release 3.4),” <https://github.com/Microsoft/SEAL>, Oct. 2019, Microsoft Research, Redmond, WA.
 - [81] J.-C. Bajard, J. Eynard, M. A. Hasan, and V. Zucca, “A full RNS variant of FV like somewhat homomorphic encryption schemes,” in *International Conference on Selected Areas in Cryptography*. Springer, 2016, pp. 423–442.
 - [82] S. Halevi, Y. Polyakov, and V. Shoup, “An improved RNS variant of the BFV homomorphic encryption scheme,” in *Cryptographers’ Track at the RSA Conference*. Springer, 2019, pp. 83–105.
 - [83] A. Q. A. Al Badawi, Y. Polyakov, K. M. M. Aung, B. Veeravalli, and K. Rohloff, “Implementation and performance evaluation of RNS variants of the BFV homomorphic encryption scheme,” *IEEE Transactions on Emerging Topics in Computing*, 2019.
 - [84] H. Chen and K. Han, “Homomorphic lower digits removal and improved FHE bootstrapping,” in *Advances in Cryptology – EUROCRYPT 2018*, J. B. Nielsen and V. Rijmen, Eds. Cham: Springer International Publishing, 2018, pp. 315–337.
 - [85] H. Chen, K. Laine, R. Player, and Y. Xia, “High-precision arithmetic in homomorphic encryption,” in *Cryptographers’ Track at the RSA Conference*. Springer, 2018, pp. 116–136.
 - [86] J. Hoffstein and J. Silverman, “Optimizations for NTRU,” *Public-Key Cryptography and Computational Number Theory, De Gruyter Proceedings in Mathematics*, pp. 77–88, 2000.
 - [87] C. Bootland, W. Castryck, I. Iliashenko, and F. Vercauteren, “Efficiently processing complex-valued data in homomorphic encryption,” *Journal of Mathematical Cryptology*, vol. 14, no. 1, pp. 55–65, 2020.
 - [88] S. Arita and S. Nakasato, “Fully homomorphic encryption for point numbers,” in *International Conference on Information Security and Cryptology*. Springer, 2016, pp. 253–270.
 - [89] C. Bonte, C. Bootland, J. W. Bos, W. Castryck, I. Iliashenko, and F. Vercauteren, “Faster homomorphic function evaluation using non-integral base encoding,” in *International Conference on Cryptographic Hardware and Embedded Systems*. Springer, 2017, pp. 579–600.
 - [90] A. Jäschke and F. Armknecht, “Accelerating homomorphic computations on rational numbers,” in *International Conference on Applied Cryptography and Network Security*. Springer, 2016, pp. 405–423.
 - [91] A. Kim, Y. Polyakov, and V. Zucca, “Revisiting homomorphic encryption schemes for finite fields,” 2021.
 - [92] A. Costache, K. Laine, and R. Player, “Evaluating the effectiveness of heuristic worst-case noise analysis in FHE,” in *European Symposium on Research in Computer Security*. Springer, 2020, pp. 546–565.
 - [93] A. Costache and N. P. Smart, “Which ring based somewhat homomorphic encryption scheme is best?” in *Cryptographers’ Track at the RSA Conference*. Springer, 2016, pp. 325–340.
 - [94] J. Hoffstein, J. Pipher, and J. H. Silverman, “NTRU: A ring-based public key cryptosystem,” in *International Algorithmic Number Theory Symposium*. Springer, 1998, pp. 267–288.
 - [95] —, “Public key cryptosystem method and apparatus,” U.S. Patent 6,081,597, (2000).
 - [96] D. Stehlé and R. Steinfeld, “Making NTRU as Secure as Worst-Case Problems over Ideal Lattices,” in *Advances in Cryptology – EUROCRYPT 2011*, K. G. Paterson, Ed. Berlin, Heidelberg: Springer, 2011, pp. 27–47.
 - [97] A. López-Alt, E. Tromer, and V. Vaikuntanathan, “On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption,” in *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, 2012, pp. 1219–1234.
 - [98] J. W. Bos, K. Lauter, J. Loftus, and M. Naehrig, “Improved security for a ring-based fully homomorphic encryption scheme,” in *IMA International Conference on Cryptography and Coding*. Springer, 2013, pp. 45–64.
 - [99] M. Albrecht, S. Bai, and L. Ducas, “A Subfield Lattice Attack on Overstretched NTRU Assumptions,” in *Advances in Cryptology – CRYPTO 2016*, M. Robshaw and J. Katz, Eds. Berlin, Heidelberg: Springer, 2016.
 - [100] J. H. Cheon, J. Jeong, and C. Lee, “An algorithm for NTRU problems and cryptanalysis of the GGH multilinear map without a low-level encoding of zero,” *LMS Journal of Computation and Mathematics*, vol. 19, no. A, pp. 255–266, 2016.
 - [101] Y. Doröz and B. Sunar, “Flattening NTRU for Evaluation Key Free Homomorphic Encryption,” *IACR Cryptology ePrint Archive*, vol. 2016, p. 315, 2016.
 - [102] T. Lepoint and M. Naehrig, “A comparison of the homomorphic encryption schemes FV and YASHE,” in *International Conference on Cryptology in Africa*. Springer, 2014, pp. 318–335.
 - [103] M. Kim and K. Lauter, “Private genome analysis through homomorphic encryption,” in *BMC medical informatics and decision making*, vol. 15, no. 5. BioMed Central, 2015, pp. 1–12.
 - [104] Z. Brakerski and V. Vaikuntanathan, “Lattice-based FHE as secure as PKE,” in *Proceedings of the 5th conference on Innovations in theoretical computer science*, 2014, pp. 1–12.
 - [105] A. Khedr, G. Gulak, and V. Vaikuntanathan, “SHIELD: scalable homomorphic implementation of encrypted data-classifiers,” *IEEE Transactions on Computers*, vol. 65, no. 9, pp. 2848–2858, 2015.
 - [106] J. Alperin-Sheriff and C. Peikert, “Faster bootstrapping with polynomial error,” in *Advances in Cryptology – CRYPTO 2014*, J. A. Garay and R. Gennaro, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 297–314.
 - [107] —, “Practical bootstrapping in quasilinear time,” in *Advances in Cryptology – CRYPTO 2013*, R. Canetti and J. A. Garay, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 1–20.
 - [108] R. Hiromasa, M. Abe, and T. Okamoto, “Packing messages and optimizing bootstrapping in GSW-FHE,” *IEICE TRANSACTIONS on Fundamentals of Electronics, Communications and Computer Sciences*, vol. 99, no. 1, pp. 73–82, 2016.

- [109] L. Ducas and D. Micciancio, “FHEW: Bootstrapping Homomorphic Encryption in Less Than a Second,” in *Advances in Cryptology – EUROCRYPT 2015*, E. Oswald and M. Fischlin, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015, pp. 617–640.
- [110] M. Frigo and S. G. Johnson, “The design and implementation of fftw3,” *Proceedings of the IEEE*, vol. 93, no. 2, pp. 216–231, 2005.
- [111] I. Chillotti, N. Gama, M. Georgieva, and M. Izabachène, “Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds,” in *international conference on the theory and application of cryptology and information security*. Springer, 2016, pp. 3–33.
- [112] —, “Faster packed homomorphic operations and efficient circuit bootstrapping for TFHE,” in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2017, pp. 377–408.
- [113] I. Chillotti, D. Ligier, J.-B. Orfila, and S. Tap, “Improved programmable bootstrapping with larger precision and efficient arithmetic circuits for tfhe,” *IACR Cryptology ePrint Archive*, 2021.
- [114] S. Carpo, N. Gama, M. Georgieva, and J. R. Troncoso-Pastoriza, “Privacy-preserving semi-parallel logistic regression training with fully homomorphic encryption,” *Cryptology ePrint Archive*, Report 2019/101, 2019, <https://eprint.iacr.org/2019/101>.
- [115] A. Guimarães, E. Borin, and D. F. Aranha, “Revisiting the functional bootstrap in TFHE,” *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 229–253, 2021.
- [116] H. Chen, I. Chillotti, and Y. Song, “Multi-key homomorphic encryption from tfhe,” in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2019, pp. 446–472.
- [117] M. Joye, “Guide to fully homomorphic encryption over the [discretized] torus,” *Cryptology ePrint Archive*, 2021.
- [118] J. H. Cheon, K. Han, A. Kim, M. Kim, and Y. Song, “Bootstrapping for approximate homomorphic encryption,” in *Advances in Cryptology – EUROCRYPT 2018*, J. B. Nielsen and V. Rijmen, Eds. Cham: Springer International Publishing, 2018, pp. 360–384.
- [119] —, “A full RNS variant of approximate homomorphic encryption,” in *International Conference on Selected Areas in Cryptography – SAC 2018*. Springer, 2018, pp. 347–368.
- [120] H. Chen, I. Chillotti, and Y. Song, “Improved bootstrapping for approximate homomorphic encryption,” in *Advances in Cryptology – EUROCRYPT 2019*, Y. Ishai and V. Rijmen, Eds. Cham: Springer International Publishing, 2019, pp. 34–54.
- [121] K. Han and D. Ki, “Better bootstrapping for approximate homomorphic encryption,” in *Topics in Cryptology – CT-RSA 2020*, S. Jarecki, Ed. Cham: Springer International Publishing, 2020, pp. 364–390.
- [122] F. Boemer, A. Costache, R. Cammarota, and C. Wierzyński, “NGraph-HE2: A High-Throughput Framework for Neural Network Inference on Encrypted Data,” in *Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, ser. WAHC ’19. New York, NY, USA: Association for Computing Machinery, 2019, p. 45–56.
- [123] D. Kim and Y. Song, “Approximate homomorphic encryption over the conjugate-invariant ring,” in *International Conference on Information Security and Cryptology*. Springer, 2018, pp. 85–102.
- [124] B. Li and D. Micciancio, “On the security of homomorphic encryption on approximate numbers,” in *Advances in Cryptology – EUROCRYPT 2021*, A. Canteaut and F.-X. Standaert, Eds. Cham: Springer International Publishing, 2021, pp. 648–677.
- [125] J. H. Cheon, S. Hong, and D. Kim, “Remark on the security of CKKS scheme in practice,” *IACR Cryptol. ePrint Arch*, vol. 2020, p. 1581, 2020.
- [126] C. Boura, N. Gama, M. Georgieva, and D. Jethchev, “CHIMERA: Combining ring-LWE-based fully homomorphic encryption schemes,” *Cryptology ePrint Archive*, Report 2018/758. <https://eprint.iacr.org/2018/758>, Tech. Rep., 2018.
- [127] “Idash privacy & security workshop 2018 - secure genome analysis competition,” <http://www.humangenomeprivacy.org/2018/index.html>, Dec. 2018.
- [128] Y. Doröz, J. Hoffstein, J. Pipher, J. H. Silverman, B. Sunar, W. Whyte, and Z. Zhang, “Fully homomorphic encryption from the finite field isomorphism problem,” in *IACR International Workshop on Public Key Cryptography*. Springer, 2018, pp. 125–155.
- [129] A. Joux, “Fully homomorphic encryption modulo fermat numbers,” *Cryptology ePrint Archive*, Report 2019/187, 2019, <https://eprint.iacr.org/2019/187>.
- [130] M. R. Albrecht, R. Player, and S. Scott, “On the concrete hardness of learning with errors,” *Journal of Mathematical Cryptology*, vol. 9, no. 3, pp. 169–203, 2015.
- [131] W. Castryck, I. Iliashenko, and F. Vercauteren, “On error distributions in ring-based LWE,” *LMS Journal of Computation and Mathematics*, vol. 19, no. A, pp. 130–145, 2016.
- [132] —, “Provably weak instances of ring-lwe revisited,” in *Advances in Cryptology – EUROCRYPT 2016*, M. Fischlin and J.-S. Coron, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 147–167.
- [133] C. Peikert, “How (not) to instantiate Ring-LWE,” in *International Conference on Security and Cryptography for Networks*. Springer, 2016, pp. 411–430.
- [134] M. R. Albrecht, “On Dual Lattice Attacks Against Small-Secret LWE and Parameter Choices in HELIB and SEAL,” in *Advances in Cryptology – EUROCRYPT 2017*, J.-S. Coron and J. B. Nielsen, Eds. Cham: Springer International Publishing, 2017, pp. 103–129.
- [135] E. Alkim, A. Roberto, B. Joppe, L. Ducas, d. I. P. Antonio, P. Schwabe, S. Douglas *et al.*, “NewHope: algorithm specifications and supporting documentation (2019),” 2019, <https://csrc.nist.gov/projects/post-quantum-cryptography/round-2-submissions>.
- [136] M. R. Albrecht and L. Ducas, “Lattice attacks on NTRU and LWE: A history of refinements,” 2021.
- [137] N. Bindel, J. Buchmann, F. Göpfert, and M. Schmidt, “Estimation of the hardness of the learning with errors problem with a restricted number of samples,” *Journal of Mathematical Cryptology*, vol. 13, no. 1, pp. 47–67, 2019.
- [138] M. R. Albrecht, “Lattice Estimator, Rebooted,” 2021, <https://martinalbrecht.wordpress.com/2021/12/21/lattice-estimator-rebooted/>.
- [139] C.-P. Schnorr and M. Euchner, “Lattice basis reduction: Improved practical algorithms and solving subset sum problems,” *Mathematical programming*, vol. 66, no. 1-3, pp. 181–199, 1994.
- [140] A. K. Lenstra, H. W. Lenstra, and L. Lovász, “Factoring polynomials with rational coefficients,” 1982.
- [141] Y. Chen and P. Q. Nguyen, “BKZ 2.0: Better lattice security estimates,” in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2011, pp. 1–20.
- [142] S. Bai, D. Stehlé, and W. Wen, “Measuring, simulating and exploiting the head concavity phenomenon in BKZ,” in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2018, pp. 369–404.
- [143] M. R. Albrecht, L. Ducas, G. Herold, E. Kirshanova, E. W. Postlethwaite, and M. Stevens, “The general sieve kernel and new records in lattice reduction,” in *Advances in Cryptology – EUROCRYPT 2019*, Y. Ishai and V. Rijmen, Eds. Cham: Springer International Publishing, 2019, pp. 717–746.
- [144] M. R. Albrecht, S. Bai, P.-A. Fouque, P. Kirchner, D. Stehlé, and W. Wen, “Faster enumeration-based lattice reduction: root hermite factor $k^{\frac{1}{2k}}$ time $k^{\frac{k}{8} + o(k)}$,” in *Advances in Cryptology – CRYPTO 2020*. Springer, 2020, pp. 186–212.
- [145] M. R. Albrecht, S. Bai, J. Li, and J. Rowell, “Lattice reduction with approximate enumeration oracles,” in *Advances in Cryptology – CRYPTO 2021*. Springer, 2021, pp. 732–759.
- [146] N. Gama and P. Q. Nguyen, “Predicting lattice reduction,” in *Advances in Cryptology – EUROCRYPT 2008*, N. Smart, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 31–51.
- [147] R. Lindner and C. Peikert, “Better key sizes (and attacks) for LWE-based encryption,” in *Cryptographers’ Track at the RSA Conference*. Springer, 2011, pp. 319–339.
- [148] M. Liu and P. Q. Nguyen, “Solving BDD by enumeration: An update,” in *Cryptographers’ Track at the RSA Conference*. Springer, 2013, pp. 293–309.
- [149] M. R. Albrecht, C. Cid, J.-C. Faugere, R. Fitzpatrick, and L. Perret, “On the complexity of the bkz algorithm on lwe,” *Designs, Codes and Cryptography*, vol. 74, no. 2, pp. 325–354, 2015.
- [150] Y. Chen, “Réduction déréseau et sécurisé concrete du chiffrement complètement homomorphe,” Ph.D. dissertation, Paris 7, 2013.
- [151] G. Hanrot, X. Pujol, and D. Stehlé, “Algorithms for the shortest and closest lattice vector problems,” in *International Conference on Coding and Cryptology*. Springer, 2011, pp. 159–190.
- [152] L. Babai, “On Lovász’ lattice reduction and the nearest lattice point problem,” *Combinatorica*, vol. 6, no. 1, pp. 1–13, 1986.
- [153] N. Gama, P. Q. Nguyen, and O. Regev, “Lattice enumeration using extreme pruning,” in *Advances in Cryptology – EUROCRYPT 2010*, H. Gilbert, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 257–278.
- [154] J. Buchmann, F. Göpfert, R. Player, and T. Wunderer, “On the hardness of lwe with binary error: Revisiting the hybrid lattice-reduction and meet-in-the-middle attack,” in *International Conference on Cryptology in Africa*. Springer, 2016, pp. 24–43.

- [155] N. Howgrave-Graham, "A Hybrid Lattice-Reduction and Meet-in-the-Middle Attack Against NTRU," in *Advances in Cryptology - CRYPTO 2007*, A. Menezes, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 150–169.
- [156] T. Wunderer, "On the security of lattice-based cryptography against lattice reduction and hybrid attacks," 2018.
- [157] S. Bai and S. D. Galbraith, "Lattice decoding attacks on binary LWE," in *Australasian Conference on Information Security and Privacy*. Springer, Cham, 2014, pp. 322–337.
- [158] R. Kannan, "Minkowski's convex body theorem and integer programming," *Mathematics of operations research*, vol. 12, no. 3, pp. 415–440, 1987.
- [159] M. R. Albrecht, R. Fitzpatrick, and F. Göpfert, "On the efficacy of solving LWE by reduction to unique-SVP," in *International Conference on Information Security and Cryptology*. Springer, 2013, pp. 293–310.
- [160] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe, "Post-Quantum Key Exchange: A New Hope," in *Proceedings of the 25th USENIX Conference on Security Symposium*, ser. SEC'16. USA: USENIX Association, 2016, p. 327–343.
- [161] M. R. Albrecht, F. Göpfert, F. Virdia, and T. Wunderer, "Revisiting the expected cost of solving uSVP and applications to LWE," in *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 2017, pp. 297–322.
- [162] S. Bai, S. Miller, and W. Wen, "A Refined Analysis of the Cost for Solving LWE via uSVP," in *International Conference on Cryptology in Africa*. Springer, 2019, pp. 181–205.
- [163] D. Dachman-Soled, L. Ducas, H. Gong, and M. Rossi, "LWE with side information: Attacks and concrete security estimation," in *Advances in Cryptology - CRYPTO 2020*, D. Micciancio and T. Ristenpart, Eds. Cham: Springer International Publishing, 2020, pp. 329–358.
- [164] H. Chen, L. Chua, K. E. Lauter, and Y. Song, "On the concrete security of lwe with small secret," *IACR Cryptol. ePrint Arch.*, vol. 2020, p. 539, 2020.
- [165] E. W. Postlethwaite and F. Virdia, "On the success probability of solving unique svp via bkz," in *Public Key Cryptography (1)*, 2021, pp. 68–98.
- [166] J. H. Cheon, M. Hhan, S. Hong, and Y. Son, "A Hybrid of Dual and Meet-in-the-Middle Attack on Sparse and Ternary Secret LWE," *IEEE Access*, vol. 7, pp. 89 497–89 506, 2019.
- [167] N. Howgrave-Graham, J. H. Silverman, and W. Whyte, "A Meet-in-the-Middle Attack on an NTRU Private key," Technical report, NTRU Cryptosystems, June 2003. Report, Tech. Rep., 2003.
- [168] T. Espitau, A. Joux, and N. Kharchenko, "On a dual/hybrid approach to small secret LWE," in *International Conference on Cryptology in India*. Springer, 2020, pp. 440–462.
- [169] S. Arora and R. Ge, "New algorithms for learning in presence of errors," in *International Colloquium on Automata, Languages, and Programming*. Springer, 2011, pp. 403–415.
- [170] Y. Elias, K. E. Lauter, E. Ozman, and K. E. Stange, "Provably weak instances of ring-lwe," in *Advances in Cryptology - CRYPTO 2015*, R. Gennaro and M. Robshaw, Eds. Berlin, Heidelberg: Springer, 2015, pp. 63–92.
- [171] H. Chen, K. Lauter, and K. E. Stange, "Attacks on the search RLWE problem with small errors," *SIAM Journal on Applied Algebra and Geometry*, vol. 1, no. 1, pp. 665–682, 2017.
- [172] —, "Security considerations for galois non-dual rlwe families," in *International Conference on Selected Areas in Cryptography*. Springer, 2016, pp. 443–462.
- [173] V. Lyubashevsky, C. Peikert, and O. Regev, "A Toolkit for Ring-LWE Cryptography," in *Advances in Cryptology - EUROCRYPT 2013*, T. Johansson and P. Q. Nguyen, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 35–54.
- [174] E. Crockett and C. Peikert, "Λoλ: Functional Lattice Cryptography," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 993–1005.
- [175] B. R. Curtis and R. Player, "On the feasibility and impact of standardising sparse-secret lwe parameter sets for homomorphic encryption," in *Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, 2019, pp. 1–10.
- [176] S. Laur, H. Lipmaa, and T. Mielikäinen, "Cryptographically private support vector machines," 08 2006, pp. 618–624.
- [177] S. Park, J. Byun, J. Lee, J. H. Cheon, and J. Lee, "He-friendly algorithm for privacy-preserving svm training," *IEEE Access*, vol. 8, pp. 57 414–57 425, 2020.
- [178] Y. Rahulamathavn BSc(Hons), PhD, R. Phan, S. Veluru, K. Cumanan, and M. Rajarajan, "Privacy-preserving multi-class support vector machine for outsourcing the data classification in cloud," *Dependable and Secure Computing, IEEE Transactions on*, vol. 11, pp. 467–479, 09 2014.
- [179] E. Makri, D. Rotaru, N. Smart, and F. Vercauteren, *EPIC: Efficient Private Image Classification (or: Learning from the Masters): The Cryptographers' Track at the RSA Conference 2019, San Francisco, CA, USA, March 4–8, 2019, Proceedings*, 01 2019, pp. 473–492.
- [180] T. Graepel, K. Lauter, and M. Naehrig, "ML confidential: Machine learning on encrypted data," vol. 7839, 11 2012, pp. 1–21.
- [181] R. Bost, R. Popa, S. Tu, and S. Goldwasser, "Machine learning classification over encrypted data," 01 2015.
- [182] L. Aslett, P. Esperança, and C. Holmes, "Encrypted statistical machine learning: new privacy preserving methods," 08 2015.
- [183] A. Khedr, G. Gulak, and V. Vaikuntanathan, "Shield: Scalable homomorphic implementation of encrypted data-classifiers," *IEEE Transactions on Computers*, vol. 65, pp. 1–1, 11 2015.
- [184] P. Li, J. Li, Z. Huang, T. Li, C.-Z. Gao, S.-M. Yiu, and K. Chen, "Multi-key privacy-preserving deep learning in cloud computing," *Future Generation Computer Systems*, vol. 74, 03 2017.
- [185] N. Dowlin, R. Gilad-Bachrach, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing, "Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy," in *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ser. ICML'16. JMLR.org, 2016, p. 201–210.
- [186] Q. Zhang, L. Yang, Z. Chen, P. Li, and M. Deen, "Privacy-preserving double-projection deep computation model with crowdsourcing on cloud for big data feature learning," *IEEE Internet of Things Journal*, vol. PP, pp. 1–1, 07 2017.
- [187] A. Brutkus, O. Elisha, and R. Gilad-Bachrach, "Low latency privacy preserving inference," 12 2018.
- [188] E. Hesamifard, D. Takabi, M. Ghasemi, and R. Wright, "Privacy-preserving machine learning as a service," *Proceedings on Privacy Enhancing Technologies*, vol. 2018, pp. 123–142, 06 2018.
- [189] A. Al Badawi, J. Chao, J. Lin, C. Mun, S. Jie, B. Tan, X. Nan, K. Aung, and V. Chandrasekhar, "The alexnet moment for homomorphic encryption: Hcnn, the first homomorphic cnn on encrypted data with gpus," 11 2018.
- [190] "Zama," <https://zama.ai>.
- [191] "Intel," <https://www.intel.com>.
- [192] "Ant group," <https://www.antgroup.com/en>.
- [193] "Better together: Privacy-preserving machine learning powered by intel sgx and intel dl boost," <https://www.intel.com/content/www/us/en/artificial-intelligence/posts/alibaba-privacy-preserving-machine-learning.html>.
- [194] "Duality technologies," <https://dualitytech.com>.
- [195] "Darpa selects researchers to accelerate use of fully homomorphic encryption," <https://www.darpa.mil/news-events/2021-03-08>.
- [196] CISCO, "The internet of things reference model," *white paper*, 2014.
- [197] —, "Fog computing and the internet of things: Extend the cloud to where the things are," *white paper*, 2015, https://www.cisco.com/c/dam/en_us/solutions/trends/iot/docs/computing-overview.pdf.
- [198] ETSI, "Multi access edge computing (mec); terminology," *ETSI Industry Specification Group*, 2019, https://www.etsi.org/deliver/etsi_gs/MEC/001_099/001/02.01.01_60/gs_MEC001v020101p.pdf.
- [199] S. P. Mohanty, U. Choppali, and E. Kougianos, "Everything you wanted to know about smart cities: The internet of things is the backbone," *IEEE Consumer Electronics Magazine*, vol. 5, no. 3, pp. 60–70, 2016.
- [200] M. Yannuzzi, F. van Lingem, A. Jain, O. L. Parellada, M. M. Flores, D. Carrera, J. L. Pérez, D. Montero, P. Chacin, A. Corsaro, and A. Olive, "A new era for cities with fog computing," *IEEE Internet Computing*, vol. 21, no. 2, pp. 54–67, 2017.
- [201] Z. Erkin, J. R. Troncoso-pastoriza, R. Legendijk, and F. Perez-Gonzalez, "Privacy-preserving data aggregation in smart metering systems: an overview," *IEEE Signal Processing Magazine*, vol. 30, no. 2, pp. 75–86, 2013.
- [202] L. Zhu, M. Li, Z. Zhang, C. Xu, R. Zhang, X. Du, and N. Guizani, "Privacy-preserving authentication and data aggregation for fog-based smart grid," *IEEE Communications Magazine*, vol. 57, no. 6, pp. 80–85, 2019.
- [203] J. H. Cheon and J. Kim, "A hybrid scheme of public-key encryption and somewhat homomorphic encryption," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 5, pp. 1052–1063, 2015.
- [204] P. Hebborn and G. Leander, *IACR Transactions on Symmetric Cryptology*, vol. 2020, no. 3, pp. 46–86, Sep. 2020.

- [205] P. Méaux, C. Carlet, A. Journault, and F.-X. Standaert, “Improved filter permutators for efficient fhe: Better instances and implementations,” in *Progress in Cryptology – INDOCRYPT 2019*, F. Hao, S. Ruj, and S. Sen Gupta, Eds. Cham: Springer International Publishing, 2019, pp. 68–91.
- [206] C. Dobraunig, L. Grassi, L. Helmlinger, C. Rechberger, M. Schafneger, and R. Walch, “Pasta: A case for hybrid homomorphic encryption,” Cryptology ePrint Archive, Report 2021/731, 2021, <https://ia.cr/2021/731>.
- [207] D. Derler, S. Ramacher, and D. Slamanig, “Homomorphic proxy re-authenticators and applications to verifiable multi-user data aggregation,” in *Financial Cryptography and Data Security*, A. Kiayias, Ed. Cham: Springer International Publishing, 2017, pp. 124–142.
- [208] Y. Kawai, T. Matsuda, T. Hirano, Y. Koseki, and G. Hanaoka, “Proxy Re-Encryption That Supports Homomorphic Operations for Re-Encrypted Ciphertexts,” *IEICE Transactions on Fundamentals of Electronics Communications and Computer Sciences*, vol. 102, no. 1, pp. 81–98, Jan. 2019.
- [209] C. Ma, J. Li, and W. Ouyang, “A homomorphic proxy re-encryption from lattices,” 11 2016, pp. 353–372.
- [210] S. Yasuda, Y. Koseki, R. Hiromasa, and Y. Kawai, “Multi-key homomorphic proxy re-encryption,” in *Information Security*, L. Chen, M. Manulis, and S. Schneider, Eds. Cham: Springer International Publishing, 2018, pp. 328–346.
- [211] Z. Li, C. Ma, and D. Wang, “Towards multi-hop homomorphic identity-based proxy re-encryption via branching program,” *IEEE Access*, vol. 5, pp. 16 214–16 228, 2017.
- [212] —, “Achieving multi-hop PRE via branching program,” *IEEE Transactions on Cloud Computing*, vol. 8, no. 1, pp. 45–58, 2020.
- [213] J. Li, Z. Qiao, K. Zhang, and C. Cui, “A lattice-based homomorphic proxy re-encryption scheme with strong anti-collusion for cloud computing,” *Sensors*, vol. 21, no. 1, 2021.
- [214] F. Luo, S. Al-Kuwari, W. Susilo, and D. H. Duong, “Chosen-ciphertext secure homomorphic proxy re-encryption,” *IEEE Transactions on Cloud Computing*, pp. 1–1, 2020.
- [215] A. Cohen, “What about Bob? the inadequacy of CPA security for proxy reencryption,” in *Public-Key Cryptography – PKC 2019*, D. Lin and K. Sako, Eds. Cham: Springer International Publishing, 2019, pp. 287–316.
- [216] R. Canetti, S. Raghuraman, S. Richelson, and V. Vaikuntanathan, “Chosen-ciphertext secure fully homomorphic encryption,” in *Public-Key Cryptography – PKC 2017*, S. Fehr, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2017, pp. 213–240.
- [217] P. K. Shantu and K. Chandrasekaran, “Generating privacy-preserved recommendation using homomorphic authenticated encryption,” in *2016 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM)*, 2016, pp. 46–53.
- [218] J. H. Cheon, K. Han, S.-M. Hong, H. J. Kim, J. Kim, S. Kim, H. Seo, H. Shim, and Y. Song, “Toward a secure drone system: Flying with real-time homomorphic authenticated encryption,” *IEEE Access*, vol. 6, pp. 24 325–24 339, 2018.
- [219] C. Joo and A. Yun, “Homomorphic authenticated encryption secure against chosen-ciphertext attack,” in *Advances in Cryptology – ASIACRYPT 2014*, P. Sarkar and T. Iwata, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 173–192.
- [220] J. Kim and A. Yun, “Secure fully homomorphic authenticated encryption,” *IEEE Access*, vol. 9, pp. 107 279–107 297, 2021.
- [221] D. Boneh, D. Freeman, J. Katz, and B. Waters, “Signing a linear subspace: Signature schemes for network coding,” in *Public Key Cryptography – PKC 2009*, S. Jarecki and G. Tsudik, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 68–87.
- [222] W. Chen, H. Lei, and K. Qi, “Lattice-based linearly homomorphic signatures in the standard model,” *Theor. Comput. Sci.*, pp. 47–54, 2016.
- [223] D. M. Freeman, “Improved security for linearly homomorphic signatures: A generic framework,” in *Public Key Cryptography – PKC 2012*, M. Fischlin, J. Buchmann, and M. Manulis, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 697–714.
- [224] B. Libert, T. Peters, M. Joye, and M. Yung, “Linearly homomorphic structure-preserving signatures and their applications,” in *Advances in Cryptology – CRYPTO 2013*, R. Canetti and J. A. Garay, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 289–307.
- [225] N. Attrapadung and B. Libert, “Homomorphic network coding signatures in the standard model,” in *Public Key Cryptography – PKC 2011*, D. Catalano, N. Fazio, R. Gennaro, and A. Nicolosi, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 17–34.
- [226] A. Esfahani, G. Mantas, and J. Rodriguez, “An efficient null space-based homomorphic mac scheme against tag pollution attacks in rlnc,” *IEEE Communications Letters*, vol. 20, no. 5, pp. 918–921, 2016.
- [227] D. Catalano, D. Fiore, and B. Warinschi, “Homomorphic signatures with efficient verification for polynomial functions?”
- [228] D. Boneh and D. M. Freeman, “Homomorphic signatures for polynomial functions,” in *Advances in Cryptology – EUROCRYPT 2011*, K. G. Paterson, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 149–168.
- [229] R. Hiromasa, Y. Manabe, and T. Okamoto, “Homomorphic signatures for polynomial functions with shorter signatures,” 2013.
- [230] S. Gorbunov, V. Vaikuntanathan, and D. Wichs, “Leveled fully homomorphic signatures from standard lattices,” in *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, ser. STOC ’15. New York, NY, USA: Association for Computing Machinery, 2015, p. 469–477.
- [231] R. Gennaro and D. Wichs, “Fully homomorphic message authenticators,” in *Advances in Cryptology – ASIACRYPT 2013*, K. Sako and P. Sarkar, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 301–320.
- [232] X. Boyen, X. Fan, and E. Shi, “Adaptively secure fully homomorphic signatures based on lattices,” Cryptology ePrint Archive, Report 2014/916, 2014.
- [233] C. Wang, B. Wu, and H. Yao, “Leveled adaptively strong-unforgeable identity-based fully homomorphic signatures,” *IEEE Access*, vol. 8, pp. 119 431–119 447, 2020.
- [234] C. G. A. Nitulescu and E. Soria-Vazquez, “Rinocchio: Snarks for ring arithmetic,” Cryptology ePrint Archive, Report 2021/322, 2021.
- [235] D. Fiore, A. Nitulescu, and D. Pointcheval, “Boosting verifiable computation on encrypted data,” Cryptology ePrint Archive, Report 2020/132, 2020.
- [236] A. Bois, I. Cascudo, D. Fiore, and D. Kim, “Flexible and efficient verifiable computation on encrypted data,” Cryptology ePrint Archive, Report 2020/1526, 2020.
- [237] E. Soria-Vázquez, “Towards secure multi-party computation on the internet: Few rounds and many parties,” Ph.D. dissertation, University of Bristol, 2019.
- [238] A. Aly and M. Van Vyve, “Practically efficient secure single-commodity multi-market auctions,” in *Financial Cryptography and Data Security*, J. Grossklags and B. Preneel, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2017, pp. 110–129.
- [239] M. A. Mustafa, S. Cleemput, A. Aly, and A. Abidin, “A secure and privacy-preserving protocol for smart metering operational data collection,” *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 6481–6490, 2019.
- [240] M. Keller, “MP-SPDZ: A versatile framework for multi-party computation,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1575–1590. [Online]. Available: <https://doi.org/10.1145/3372297.3417872>
- [241] M. Keller, V. Pastro, and D. Rotaru, “Overdrive: Making spdz great again,” in *Advances in Cryptology – EUROCRYPT 2018*, J. B. Nielsen and V. Rijmen, Eds. Cham: Springer International Publishing, 2018, pp. 158–189.
- [242] A. Aly, E. Orsini, D. Rotaru, N. P. Smart, and T. Wood, “Zaphod: Efficiently Combining LSSS and Garbled Circuits in SCALE,” in *Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, ser. WAHC ’19. New York, NY, USA: Association for Computing Machinery, 2019, p. 33–44.
- [243] Y. Lindell, N. P. Smart, and E. Soria-Vazquez, “More efficient constant-round multi-party computation from bmr and she,” in *Theory of Cryptography*, M. Hirt and A. Smith, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 554–581.
- [244] H. Chen, M. Kim, I. Razenshteyn, D. Rotaru, Y. Song, and S. Wagh, “Maliciously secure matrix multiplication with applications to private deep learning,” in *Advances in Cryptology – ASIACRYPT 2020*, S. Moriai and H. Wang, Eds. Cham: Springer International Publishing, 2020, pp. 31–59.
- [245] M. Keller, E. Orsini, and P. Scholl, “Mascot: Faster malicious arithmetic secure computation with oblivious transfer,” Cryptology ePrint Archive, Report 2016/505, 2016.
- [246] C. Hazay, P. Scholl, and E. Soria-Vazquez, “Low cost constant round mpc combining bmr and oblivious transfer,” in *Advances in Cryptology – ASIACRYPT 2017*, T. Takagi and T. Peyrin, Eds. Cham: Springer International Publishing, 2017, pp. 598–628.
- [247] P. Mukherjee and D. Wichs, “Two round multiparty computation via multi-key fhe,” in *Advances in Cryptology – EUROCRYPT 2016*,

- M. Fischlin and J.-S. Coron, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 735–763.
- [248] E. Kim, H.-S. Lee, and J. Park, “Towards round-optimal secure multiparty computations: Multikey fhe without a crs,” *International Journal of Foundations of Computer Science*, vol. 31, no. 02, pp. 157–174, 2020.
- [249] PALISADE, <https://palisade-crypto.org>.
- [250] Lattigo, <http://github.com/ldsec/lattigo>.
- [251] FHEW, <https://github.com/lducas/FHEW>.
- [252] TFHE, <https://github.com/tfhe/tfhe>.
- [253] Concrete, <https://github.com/zama-ai/concrete>.
- [254] HEAAN, <https://github.com/snucrypto/HEAAN>.
- [255] RNS-HEAAN, <https://github.com/KyoohyungHan/FullRNS-HEAAN>.
- [256] FV-NFLlib, <https://github.com/CryptoExperts/FV-NFLlib>.
- [257] cuFHE, <https://github.com/vernamlab/cuFHE>.
- [258] nuFHE, <https://github.com/nucypher/nuFHE>.
- [259] S. Halevi and V. Shoup, “Algorithms in helib,” in *Advances in Cryptology – CRYPTO 2014*, J. A. Garay and R. Gennaro, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 554–571.
- [260] V. Shoup, “NTL: A Library for doing Number Theory,” 2016, [Online] <https://libntl.org/>.
- [261] F. Boemer, S. Kim, G. Seifu, F. D. de Souza, V. Gopal *et al.*, “Intel HEXL (release 1.2),” <https://github.com/intel/hexl>, 2021.
- [262] C. Mouchet, J.-P. Bossuat, J. Troncoso-Pastoriza, and J.-P. Hubaux, “Lattigo: A multiparty homomorphic encryption library in go,” in *Proceedings of the 8th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, ser. WAHC ’20, 2020.
- [263] I. Chillotti, M. Joye, D. Ligier, J.-B. Orfila, and S. Tap, “CONCRETE: Concrete operates on ciphertexts rapidly by extending THe,” in *Proceedings of the 8th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, ser. WAHC ’20, vol. 15, 2020.
- [264] NFLlib, <https://github.com/quarkslab/NFLlib>.
- [265] W. Dai and B. Sunar, “cuhe: A homomorphic encryption accelerator library,” in *International Conference on Cryptography and Information Security in the Balkans*. Springer, 2015, pp. 169–186.
- [266] A. Viand, P. Jattke, and A. Hithnawi, “SoK: Fully Homomorphic Encryption Compilers,” in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2021, pp. 1092–1108.
- [267] ALCHEMY, <https://github.com/cpeikert/ALCHEMY>.
- [268] Cingulata, <https://github.com/CEA-LIST/Cingulata>.
- [269] Encrypt-Everything-Everywhere, <https://github.com/momalab/e3>.
- [270] EVA, <https://github.com/microsoft/EVA>.
- [271] Marble, <https://github.com/MarbleHE/Marble>.
- [272] D. W. Archer, J. M. Calderón Trilla, J. Dagit, A. Malozemoff, Y. Polyakov, K. Rohloff, and G. Ryan, “RAMPARTS: A programmer-friendly system for building homomorphic encryption applications,” in *Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, 2019, pp. 57–68.
- [273] Transpiler, <https://github.com/google/fully-homomorphic-encryption>.
- [274] R. Dathathri, O. Saarikivi, H. Chen, K. Laine, K. Lauter, S. Maleki, M. Musuvathi, and T. Mytkowicz, “CHET: an optimizing compiler for fully-homomorphic neural-network inferring,” in *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2019, pp. 142–156.
- [275] nGraph-HE2, <https://github.com/IntelAI/he-transformer>.
- [276] T. van Elsloo, G. Patrini, and H. Ivey-Law, “SEALion: A framework for neural network inference on encrypted data,” *arXiv preprint arXiv:1904.12840*, 2019.
- [277] E. Crockett, “Simply safe lattice cryptography,” Ph.D. dissertation, Georgia Institute of Technology, 2017.
- [278] E. Crockett, C. Peikert, and C. Sharp, “Alchemy: A language and compiler for homomorphic encryption made easy,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, 2018, pp. 1020–1037.
- [279] “ $\lambda \circ \lambda$,” <https://github.com/cpeikert/Lol>.
- [280] S. Carpov, P. Dubrulle, and R. Sirdey, “Armadillo: a compilation chain for privacy preserving applications,” in *Proceedings of the 3rd International Workshop on Security in Cloud Computing*, 2015, pp. 13–19.
- [281] D. Harvey and W. Hart, “FLINT: Fast Library for Number Theory,” 2007, [Online] <https://www.flintlib.org/index.html>.
- [282] The Sage Developers, *SageMath, the Sage Mathematics Software System*, <https://www.sagemath.org>.
- [283] E. Chielle, O. Mazonka, N. G. Tsoutsos, and M. Maniatakos, “E3: A Framework for Compiling C++ Programs with Encrypted Operands,” *IACR Cryptology ePrint Archive*, vol. 2018, p. 1013, 2018, [Online] <https://eprint.iacr.org/2018/1013.pdf>.
- [284] R. Dathathri, B. Kostova, O. Saarikivi, W. Dai, K. Laine, and M. Musuvathi, “EVA: An encrypted vector arithmetic language and compiler for efficient homomorphic computation,” in *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2020, pp. 546–561.
- [285] A. Viand and H. Shafagh, “Marble: Making fully homomorphic encryption accessible to all,” in *Proceedings of the 6th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, 2018, pp. 49–60.
- [286] S. Gorantala, R. Springer, S. Purser-Haskell, W. Lam, R. J. Wilson, A. Ali, E. P. Astor, I. Zukerman, S. Ruth, C. Dibak, P. Schoppmann, S. Kulankhina, A. Forget, D. Marn, C. Tew, R. Misoczki, B. Guillen, X. Ye, D. Kraft, D. Desfontaines, A. Krishnamurthy, M. Guevara, I. M. Perera, Y. Sushko, and B. Gipson, “A General Purpose Transpiler for Fully Homomorphic Encryption,” *arXiv preprint arXiv:2106.07893*, 2021. [Online]. Available: <https://arxiv.org/abs/2106.07893>
- [287] F. Boemer, Y. Lao, R. Cammarota, and C. Wierzynski, “nGraph-HE: a graph compiler for deep learning on homomorphically encrypted data,” in *Proceedings of the 16th ACM International Conference on Computing Frontiers*, 2019, pp. 3–13.
- [288] S. Cyphers, A. K. Bansal, A. Bhiwandiwala, J. Bobba, M. Brookhart, A. Chakraborty, W. Constable, C. Convey, L. Cook, O. Kanawi *et al.*, “Intel nGraph: An intermediate representation, compiler, and executor for deep learning,” *arXiv preprint arXiv:1801.08058*, 2018.
- [289] F. Boemer, A. Costache, R. Cammarota, and C. Wierzynski, “nGraph-HE2: A high-throughput framework for neural network inference on encrypted data,” in *Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, 2019, pp. 45–56.
- [290] Y. Doröz, E. Öztürk, and B. Sunar, “Accelerating fully homomorphic encryption in hardware,” *IEEE Transactions on Computers*, vol. 64, no. 6, pp. 1509–1521, 2014.
- [291] D. B. Cousins, K. Rohloff, and D. Sumorok, “Designing an FPGA-accelerated homomorphic encryption co-processor,” *IEEE Transactions on Emerging Topics in Computing*, vol. 5, no. 2, pp. 193–206, 2016.
- [292] S. S. Roy, F. Vercauteren, N. Mentens, D. D. Chen, and I. Verbauwhede, “Compact ring-LWE cryptoprocessor,” in *International workshop on cryptographic hardware and embedded systems*. Springer, 2014, pp. 371–391.
- [293] S. S. Roy, F. Turan, K. Jarvinen, F. Vercauteren, and I. Verbauwhede, “FPGA-based high-performance parallel architecture for homomorphic computing on encrypted data,” in *2019 IEEE International symposium on high performance computer architecture (HPCA)*. IEEE, 2019, pp. 387–398.
- [294] M. S. Riaz, K. Laine, B. Pelton, and W. Dai, “HEAX: An architecture for computing on encrypted data,” in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020, pp. 1295–1309.
- [295] F. Turan, S. S. Roy, and I. Verbauwhede, “HEAWS: An Accelerator for Homomorphic Encryption on the Amazon AWS FPGA,” *IEEE Transactions on Computers*, vol. 69, no. 8, pp. 1185–1196, 2020.
- [296] C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan, “GAZELLE: A low latency framework for secure neural network inference,” in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 1651–1669.
- [297] B. Reagen, W. Choi, Y. Ko, V. Lee, G.-Y. Wei, H.-H. S. Lee, and D. Brooks, “Cheetah: Optimizing and accelerating homomorphic encryption for private inference,” in *Proceedings of the 27th IEEE international symposium on High Performance Computer Architecture (HPCA-27)*, 2021.
- [298] A. Feldmann, N. Samardzic, A. Krastev, S. Devadas, R. Dreslinski, K. Eldefrawy, N. Genise, C. Peikert, and D. Sanchez, “F1: A Fast and Programmable Accelerator for Fully Homomorphic Encryption (Extended Version),” *arXiv preprint arXiv:2109.05371*, 2021.
- [299] “Cryptolab,” <https://www.cryptolab.co.kr/>.
- [300] “Protect your accounts from data breaches with Password Checkup,” 2019, [Online] <https://security.googleblog.com/2019/02/protect-your-accounts-from-data.html>.

BIOGRAPHIES



Chiara Marcolla is a senior researcher at the Cryptography Research Centre at the Technology Innovation Institute (TII). She specialized in the theory of affine-variety codes, using geometrical and algebraic methods such as the properties of Norm-Trace curves and Gröbner bases theory of zero-dimensional ideals. Her current research interests include rank-based post-quantum cryptography, searchable and fully homomorphic encryption. She received her Ph.D. degrees in Mathematics from the University of Trento (Italy) in 2013. Between 2013

and 2014, she was a postdoctoral researcher and she led the Cloud Security Team at the Department of Mathematics of the University of Trento. She was a postdoctoral researcher and had a research fellowship from the Department of Mathematics of the University of Turin (2015-2016).



Victor Sucasas received his Ph.D. on Electrical and Electronic Engineering from the 5G Innovation Centre at the University of Surrey (UK) in 2016. He was also a Marie Curie ESR at Instituto de Telecomunicações (Portugal) and a visiting researcher at University of Trento (Italy). In 2016 he became a postdoctoral researcher in network security and privacy preserving systems at Instituto de Telecomunicações where he participated in European projects ECSEL-SWARMs, CATRENE-H2O and ECSEL-SECREDAS. His research inter-

ests cover privacy-preserving authentication mechanisms, secure protocols, secure multi-party computation and privacy-preserving machine learning. He has also worked on wireless ad hoc networks, underwater networks, secure cognitive radio systems and physical layer security for wireless communications. Currently, he leads the Confidential Computing team at the Cryptography Research Centre at the Technology Innovation Institute (TII), Abu Dhabi, EAU. He is an IEEE and ComSoc member and an EAI Fellow.



Marc Manzano leads the Quantum Security group at Sandbox Quantum. Prior to that, he was a Senior Staff Software Engineer at Sandbox@Alphabet, a group within Google X devoted to quantum technologies and AI, where he focuses on research and development of quantum-secure communication solutions. His current research interests include post-quantum cryptography, lightweight cryptography, fully-homomorphic encryption, the intersection between machine learning and cryptanalysis, performance optimisations of cryptographic implementa-

tions on a wide range of architectures, and quantum algorithms. Over the past ten years, Dr Manzano has led the development of many secure cryptographic libraries and protocols. Dr Manzano was formerly the Vice President of the Cryptography Research Centre at the Technology Innovation Institute, a UAE-based scientific research centre. Prior to that, he held several positions where he was responsible for implementing pivotal cryptographic components of a variety of secure communication products, including an electronic voting platform. Dr Manzano holds a PhD in Computers Network Security, which he earned under the supervision of the University of Girona (Spain) and Kansas State University (United States). He earned an MSc in Computer Science from the University of Girona (Spain), while he did research stays at UC3M (Spain) and at DTU (Denmark). He initiated his research career while finalizing his BSc in Computer Engineering at Strathclyde University (UK).



Riccardo Bassoli is a senior researcher at the Deutsche Telekom Chair of Communication Networks, Faculty of Electrical and Computer Engineering, Technische Universität Dresden (Germany). He received his B.Sc. and M.Sc. degrees in Telecommunications Engineering from the University of Modena and Reggio Emilia (Italy) in 2008 and 2010 respectively. Next, he received his Ph.D. degree from the 5G Innovation Centre at the University of Surrey (UK), in 2016. He was also a Marie Curie ESR at the Instituto de Telecomunicações (Portugal) and visiting researcher at Airbus Defence and Space (France). Between 2016 and 2019, he was a postdoctoral researcher at the University of Trento (Italy). He is an IEEE and ComSoc member. He is also a member of Glue Technologies for Space Systems Technical Panel of IEEE AESS.



Frank H.P. Fitzek is a Professor and head of the Deutsche Telekom Chair of Communication Networks at the Technische Universität Dresden, coordinating the 5G Lab Germany. He is the spokesman of the DFG Cluster of Excellence CeTI. He received his diploma (Dipl.-Ing.) degree in electrical engineering from the University of Technology Rheinisch-Westfälische Technische Hochschule (RWTH) Aachen, Germany, in 1997 and his Ph.D. (Dr.-Ing.) in electrical engineering from the Technical University Berlin, Germany in 2002 and became

Adjunct Professor at the University of Ferrara, Italy in the same year. In 2003 he joined Aalborg University as Associate Professor and later became Professor. In 2005 he won the YRP award for the work on MIMO MDC and received the Young Elite Researcher Award of Denmark. He was selected to receive the NOKIA Champion Award several times in a row from 2007 to 2011. In 2008 he was awarded the Nokia Achievement Award for his work on cooperative networks. In 2011 he received the SAPERE AUDE research grant from the Danish government and in 2012 he received the Vodafone Innovation prize. In 2015 he was awarded the honorary degree Doctor Honoris Causa from Budapest University of Technology and Economics (BUTE).



Najwa Aaraj is the Chief Researcher at the Cryptography Research Centre at the Technology Innovation Institute (TII). Her interests include post-quantum cryptography (PQC), lightweight cryptography, hardware cryptographic cores, privacy-preserving protocols, and applied machine learning for cryptographic technologies. She is also Acting Chief Researcher at TII's Autonomous Robotics Research Centre, which is dedicated to breakthrough developments in robotics and autonomy. Dr. Aaraj earned a PhD with Highest Distinction in Applied

Cryptography and Embedded Systems Security from Princeton University (USA). She has extensive expertise in applied cryptography, trusted platforms, security architecture for embedded systems, software exploit detection and prevention systems, and biometrics. She has over 15 years of experience with global firms, working in multiple geographies from Australia to the United States. Prior to TII, Dr Aaraj assumed positions at DarkMatter, Booz & Company, IBM T.J. Watson Security Research, NY, Intel Research, Oregon, and NEC Laboratories, Princeton. Dr Aaraj has written multiple conference and journal papers, and received patents on applied cryptography, embedded system security, and machine learning-based protection of Internet of Things (IoT) systems. Dr Aaraj is on the advisory board of New York-based NeuTigers, a leading-edge startup revolutionising the next generation of energy/latency-efficient artificial intelligence (AI). She is also Adviser within the Strategic Advisory Group at Paladin Capital Group (Cyber Venture Capital) and Adjunct Professor at the Mohamed Bin Zayed Artificial Intelligence University (Machine Learning Research Group). In addition, she is an Adviser to multiple security and Machine Learning startups including Okinawa Institute of Science and Technology Graduate University.