

Challenges and open problems in Fully Homomorphic Encryption

Anamaria Costache

27 July 2022




Norwegian University of
Science and Technology

Table of contents

1. Introduction, history and applications
2. Schemes
3. Packing
4. Noise in HE
5. Evaluating functions homomorphically
6. Interlude: Iterative Algorithms and Approximate Encryption
7. Real-life implementation
8. HE as a Primitive

Introduction, history and applications

What is Homomorphic Encryption?



x Encryption of x

$F(x)$ Encryption of $F(x)$

Thanks to Rachel Player for this slide!

Fully Homomorphic Encryption: a brief history

- 1978 Proposed as open problem [RAD78]
- 2009 Gentry's blueprint [Gentry09]
- 2010 First Implementation
- 2010 [DGHV10]
- 2012 Second generation schemes: BGV [BGV12], B/FV [Bra12, FV12], ...
- 2013 Third generation schemes: GSW [GSW13], ...
First open source libraries available
- 2014 FHEW [DM14]
- 2016 TFHE introduces fast bootstrapping (less than 0.1 seconds!) [CGGI16]
- 2017 Standardisation effort begins
Approximate homomorphic encryption: CKKS [CKKS17]
- 2020 ISO efforts begin

Community efforts

Actively maintained libraries

Name	Implements	Available at
HElib	BGV, CKKS	github.com/homenc/HElib
SEAL	BFV, CKKS	github.com/Microsoft/SEAL
Palisade	BFV, BGV and more!	git.njit.edu/palisade/PALISADE
Concrete	TFHE	github.com/zama-ai/concrete
Lattigo	CKKS, BFV	github.com/tuneinsight/lattigo


Standardisation effort¹

- Homomorphic Encryption Standard recommends secure parameters
- White papers on applications, APIs, schemes etc.
- Usability focus: use-cases
- Standardisation: ISO/ ISE

¹homomorphicencryption.org, fhe.org

Many possible applications

- Medical
- Genomics
- Machine Learning
- Statistics
- Smart cities
- Cyber physical systems
- Private information retrieval
- Database search
- Private set intersection
- Electronic voting
- ...



Schemes

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x] / (\Phi_n, q)$

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x]/(\Phi_n, q)$
- Ciphertext moduli q_k

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x]/(\Phi_n, q)$
- Ciphertext moduli q_k
 - There are L primes p_0, \dots, p_{L-1}

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x] / (\Phi_n, q)$
- Ciphertext moduli q_k
 - There are L primes p_0, \dots, p_{L-1}
 - The chain of moduli is formed as $q_k = \prod_{j=0}^k p_j$

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x]/(\Phi_n, q)$
- Ciphertext moduli q_k
 - There are L primes p_0, \dots, p_{L-1}
 - The chain of moduli is formed as $q_k = \prod_{j=0}^k p_j$
 - $Q = q_{L-1} = \prod_{j=0}^{L-1} p_j$ the product all the primes is the “top” modulus

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x]/(\Phi_n, q)$
- Ciphertext moduli q_k
 - There are L primes p_0, \dots, p_{L-1}
 - The chain of moduli is formed as $q_k = \prod_{j=0}^k p_j$
 - $Q = q_{L-1} = \prod_{j=0}^{L-1} p_j$ the product all the primes is the “top” modulus
- Gaussian error distribution χ with standard deviation σ

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x]/(\Phi_n, q)$
- Ciphertext moduli q_k
 - There are L primes p_0, \dots, p_{L-1}
 - The chain of moduli is formed as $q_k = \prod_{j=0}^k p_j$
 - $Q = q_{L-1} = \prod_{j=0}^{L-1} p_j$ the product all the primes is the “top” modulus
- Gaussian error distribution χ with standard deviation σ
- Secret distribution S

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x]/(\Phi_n, q)$
- Ciphertext moduli q_k
 - There are L primes p_0, \dots, p_{L-1}
 - The chain of moduli is formed as $q_k = \prod_{j=0}^k p_j$
 - $Q = q_{L-1} = \prod_{j=0}^{L-1} p_j$ the product all the primes is the “top” modulus
- Gaussian error distribution χ with standard deviation σ
- Secret distribution S
- Plaintext modulus t^1

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x]/(\Phi_n, q)$
- Ciphertext moduli q_k
 - There are L primes p_0, \dots, p_{L-1}
 - The chain of moduli is formed as $q_k = \prod_{j=0}^k p_j$
 - $Q = q_{L-1} = \prod_{j=0}^{L-1} p_j$ the product all the primes is the “top” modulus
- Gaussian error distribution χ with standard deviation σ
- Secret distribution S
- Plaintext modulus t^1
- Security parameter λ

¹For BGV and BFV

A Ring-LWE-based (levelled) FHE scheme is parameterised by:

- Dimension n (here always a power of two)
- Power-of-two cyclotomic polynomial $\Phi_n = x^n + 1$
- Polynomial ring $R_q = \mathbb{Z}[x]/(\Phi_n, q)$
- Ciphertext moduli q_k
 - There are L primes p_0, \dots, p_{L-1}
 - The chain of moduli is formed as $q_k = \prod_{j=0}^k p_j$
 - $Q = q_{L-1} = \prod_{j=0}^{L-1} p_j$ the product all the primes is the “top” modulus
- Gaussian error distribution χ with standard deviation σ
- Secret distribution S
- Plaintext modulus t^1
- Security parameter λ
- \mathbb{Z}_q means $\mathbb{Z}/q\mathbb{Z}$, $[\cdot]_q$ means reduction modulo q

¹For BGV and BFV

The CKKS scheme - basic operations

Ciphertexts are elements $(c_0, c_1) \in R_q^2$.

SecretKeyGen(λ): Sample $s \leftarrow S$ and output $sk = (1, s)$.

The CKKS scheme - basic operations

Ciphertexts are elements $(c_0, c_1) \in R_q^2$.

SecretKeyGen(λ): Sample $s \leftarrow S$ and output $\text{sk} = (1, s)$.

PublicKeyGen(sk): Recall $\text{sk} = (1, s)$ and sample $a \leftarrow R_q$ uniformly at random and $e \leftarrow \chi$. Output $\text{pk} = ([-as + e]_q, a)$.

The CKKS scheme - basic operations

Ciphertexts are elements $(c_0, c_1) \in R_q^2$.

SecretKeyGen(λ): Sample $s \leftarrow S$ and output $\text{sk} = (1, s)$.

PublicKeyGen(sk): Recall $\text{sk} = (1, s)$ and sample $a \leftarrow R_q$ uniformly at random and $e \leftarrow \chi$. Output $\text{pk} = ([-as + e]_q, a)$.

Encrypt(pk, m): For the message $m \in R$. Let $\text{pk} = (p_0, p_1)$, sample $v \leftarrow S$ and $e_1, e_2 \leftarrow \chi$. Output $\text{ct} = ([m + e_1 + p_0v]_q, [p_1v + e_2]_q)$.

The CKKS scheme - basic operations

Ciphertexts are elements $(c_0, c_1) \in R_q^2$.

SecretKeyGen(λ): Sample $s \leftarrow S$ and output $\text{sk} = (1, s)$.

PublicKeyGen(sk): Recall $\text{sk} = (1, s)$ and sample $a \leftarrow R_q$ uniformly at random and $e \leftarrow \chi$. Output $\text{pk} = ([-as + e]_q, a)$.

Encrypt(pk, m): For the message $m \in R$. Let $\text{pk} = (p_0, p_1)$, sample $v \leftarrow S$ and $e_1, e_2 \leftarrow \chi$. Output $\text{ct} = ([m + e_1 + p_0v]_q, [p_1v + e_2]_q)$.

Decrypt(sk, ct): Let $\text{ct} = (c_0, c_1)$. Output $m' = [c_0 + c_1s]_q$.

The CKKS scheme - homomorphic operations

- Addition is coordinate-wise
- Multiplication, at a high level, corresponds to three steps:
 - **Pre-Multiply:** from two ciphertexts in R_q^2 that both decrypt under sk , we have one ciphertext in R_q^3 that decrypts under sk^2

The CKKS scheme - homomorphic operations

- Addition is coordinate-wise
- Multiplication, at a high level, corresponds to three steps:
 - Pre-Multiply: from **two** ciphertexts in R_q^2 that both decrypt under **sk**, we have **one** ciphertext in R_q^3 that decrypts under **sk²**
 - Relin: From **one** ciphertext in R_q^3 that decrypts under **sk**, we go to **one** ciphertext in R_q^2 that decrypts under **sk²**

The CKKS scheme - homomorphic operations

- Addition is coordinate-wise
- Multiplication, at a high level, corresponds to three steps:
 - Pre-Multiply: from **two** ciphertexts in R_q^2 that both decrypt under **sk**, we have **one** ciphertext in R_q^3 that decrypts under **sk²**
 - Relin: From one ciphertext in R_q^3 that decrypts under **sk**, we go to one ciphertext in R_q^2 that decrypts under **sk²**
 - Rescale: Multiply by q'/q and round to reduce the noise and output one ciphertext in $R_{q'}$ that decrypts under **sk**

The CKKS scheme - homomorphic operations

We have $ct_0 = (ct_0[0], ct_0[1])$ and $ct_1 = (ct_1[0], ct_1[1]) \pmod{q}$.

Add(ct_0, ct_1): Output $ct = ([ct_0[0] + ct_1[0]]_q, [ct_0[1] + ct_1[1]]_q)$.

The CKKS scheme - homomorphic operations

We have $ct_0 = (ct_0[0], ct_0[1])$ and $ct_1 = (ct_1[0], ct_1[1]) \pmod{q}$.

Add(ct_0, ct_1): Output $ct = ([ct_0[0] + ct_1[0]]_q, [ct_0[1] + ct_1[1]]_q)$.

Pre-Multiply(ct_0, ct_1): Set $d_0 = [ct_0[0]ct_1[0]]_q$,
 $d_1 = [ct_0[0]ct_1[1] + ct_0[1]ct_1[0]]_q$, and
 $d_2 = [ct_0[1]ct_1[1]]_q$. Output $ct = (d_0, d_1, d_2)$.

The CKKS scheme - homomorphic operations

We have $ct_0 = (ct_0[0], ct_0[1])$ and $ct_1 = (ct_1[0], ct_1[1]) \pmod{q}$.

Add(ct_0, ct_1): Output $ct = ([ct_0[0] + ct_1[0]]_q, [ct_0[1] + ct_1[1]]_q)$.

Pre-Multiply(ct_0, ct_1): Set $d_0 = [ct_0[0]ct_1[0]]_q$,
 $d_1 = [ct_0[0]ct_1[1] + ct_0[1]ct_1[0]]_q$, and
 $d_2 = [ct_0[1]ct_1[1]]_q$. Output $ct = (d_0, d_1, d_2)$.

KeySwitch(ct, evk): Let $ct[0] = d_0$, $ct[1] = d_1$ and $ct[2] = d_2$. Recall
 $evk[0] = -a's + e' + Ps^2$ and $evk[1] = a'$. Set
 $c'_0 = [d_0 + \lfloor P^{-1} \cdot d_2 \cdot (-a's + e' + Ps^2) \rfloor]_q$, and
 $c'_1 = [d_1 + \lfloor P^{-1} \cdot d_2 \cdot a' \rfloor]_q$. Output $ct' = (c'_0, c'_1)$.

The CKKS scheme - homomorphic operations

We have $ct_0 = (ct_0[0], ct_0[1])$ and $ct_1 = (ct_1[0], ct_1[1]) \pmod{q}$.

Add(ct_0, ct_1): Output $ct = ([ct_0[0] + ct_1[0]]_q, [ct_0[1] + ct_1[1]]_q)$.

Pre-Multiply(ct_0, ct_1): Set $d_0 = [ct_0[0]ct_1[0]]_q$,
 $d_1 = [ct_0[0]ct_1[1] + ct_0[1]ct_1[0]]_q$, and
 $d_2 = [ct_0[1]ct_1[1]]_q$. Output $ct = (d_0, d_1, d_2)$.

KeySwitch(ct, evk): Let $ct[0] = d_0$, $ct[1] = d_1$ and $ct[2] = d_2$. Recall
 $evk[0] = -a's + e' + Ps^2$ and $evk[1] = a'$. Set
 $c'_0 = [d_0 + [P^{-1} \cdot d_2 \cdot (-a's + e' + Ps^2)]]_q$, and
 $c'_1 = [d_1 + [P^{-1} \cdot d_2 \cdot a']]_q$. Output $ct' = (c'_0, c'_1)$.

Rescale(ct, q, q'): Let $ct = (c_0, c_1)$. Set $c'_0 = \left[\left[\frac{q'c_0}{q} \right] \right]_{q'}$, and
 $c'_1 = \left[\left[\frac{q'c_1}{q} \right] \right]_{q'}$. Output $ct = (c'_0, c'_1)$.

The BGV scheme

SecretKeyGen(λ): Sample $s \leftarrow S$ and output $\text{sk} = s$.

PublicKeyGen(sk): Set $s = \text{sk}$ and sample $a \leftarrow R_q$ uniformly at random and $e \leftarrow \chi$. Output $\text{pk} = ([-(as + et)]_q, a)$.

Encrypt(pk, m): For the message $m \in R_t$. Let $\text{pk} = (p_0, p_1)$, sample $u \leftarrow S$ and $e_1, e_2 \leftarrow \chi$. Output $\text{ct} = ([m + p_0u + te_1]_q, [p_1u + te_2]_q)$.

Decrypt(sk, ct): Let $s = \text{sk}$ and $\text{ct} = (c_0, c_1)$. Output $m' = [[c_0 + c_1s]_q]_t$.

The BFV scheme

SecretKeyGen(λ): Sample $s \leftarrow S$ and output $\text{sk} = s$.

PublicKeyGen(sk): Set $s = \text{sk}$ and sample $a \leftarrow R_q$ uniformly at random and $e \leftarrow \chi$. Output $\text{pk} = ([-(as + e)]_q, a)$.

Encrypt(pk, m): For the message $m \in R_t$. Let $\text{pk} = (p_0, p_1)$, sample $u \leftarrow S$ and $e_1, e_2 \leftarrow \chi$. Output $\text{ct} = ([\Delta m + p_0 u + e_1]_q, [p_1 u + e_2]_q)$.

Decrypt(sk, ct): Let $s = \text{sk}$ and $\text{ct} = (c_0, c_1)$. Output $m' = \left[\left[\frac{t}{q} [c_0 + c_1 s]_q \right] \right]_t$.

A side-by-side comparison

Scheme	BGV	BFV	CKKS
Message encoding	$m + t \cdot e$	$\Delta \cdot m + e$	$m + e$
Message encoding	Lower bits	Upper bits	Approximate encryption
Decryption	$m' = \lfloor [c_0 + c_1 s]_q \rfloor_t$	$m' = \left\lfloor \left\lfloor \frac{t}{q} [c_0 + c_1 s]_q \right\rfloor \right\rfloor_t$	$m' = \lfloor [c_0 + c_1 s]_q \rfloor$
Multiplication	$m_0 m_1 + t^2 e_0 e_1 + t(e_0 m_1 + e_1 m_0)$	$\Delta^2 m_0 m_1 + \Delta(e_0 m_1 + e_1 m_0) + e_0 e_1$	$m_0 m_1 + m_1 e_0 + m_0 e_1 + e_0 e_1$

Noise growth is much slower in CKKS.

Packing

Encoding and Decoding for CKKS

The CKKS scheme uses the canonical embedding to define an encoding from the message space $\mathbb{C}^{N/2}$ to the “plaintext” space $\mathbb{Z}[X]/(X^N + 1)$ in the following way: an isomorphism $\tau : \mathbb{R}[X]/(X^N + 1) \rightarrow \mathbb{C}^{N/2}$ can be defined via considering the canonical embedding restricted to $N/2$ of the $2N^{\text{th}}$ primitive roots and discarding conjugates. Encoding and decoding then use this map τ , as well as a precision parameter Δ , as follows:

$$\text{Encode}(\mathbf{z}, \Delta) = \lfloor \Delta \tau^{-1}(\mathbf{z}) \rfloor, \quad \text{Decode}(m, \Delta) = \frac{1}{\Delta} \tau(m),$$

where $\mathbf{z} \in \mathbb{C}^{N/2}$, $m \in \mathbb{Z}_q[X]/(X^N + 1)$ and $\lfloor \cdot \rfloor$ is taken coefficient-wise.

Can pack up to $N/2$ values into a single ciphertext.

Plaintext packing in BGV - CRT for rings

Theorem

Let R be a ring, and I_0, \dots, I_{k-1} be pairwise coprime ideals. Let $I = \bigcap_{i=0}^{k-1} I_i$. Then we have that

$$R/I \cong R/I_0 \times \dots \times R/I_{k-1}.$$


Now let $R = \mathbb{Z}_q[x]/(\Phi_n)$, where Φ_n is the n th cyclotomic polynomial, q is a ciphertext modulus and $\mathbb{Z}_t[x]/(\Phi_n)$ is the plaintext space. Consider $\mathbb{Z}/t\mathbb{Z}$ such that it contains a primitive n th root of unity. Then, Φ_n factors into ℓ irreducible factors

$$\Phi_n = \prod_{i=0}^{\ell-1} f_i(x) \pmod{t}.$$

The ideals (f_i) are pairwise coprime in R , and we can apply the Chinese Remainder Theorem for rings to obtain plaintext slots.


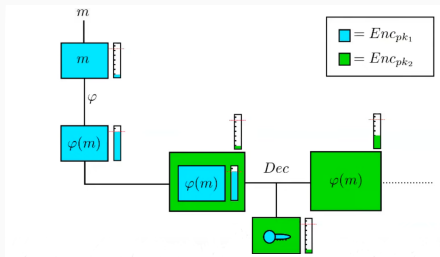
Noise in HE

The noise problem, illustrated



Noise management

Two main techniques for managing the noise: **bootstrapping** and **modulus switching** (or rescaling).



What is noise and why is it important?

Noise in homomorphic encryption

- All ciphertexts have inherent **noise**
- Noise grows during homomorphic operations
- If noise too large, decryption will fail

Good understanding of noise growth is essential

- Either need to determine when to bootstrap
- Or need to determine when to do modulus switching
- This enables us to **choose appropriate parameters**, ideally small ones
- Requiring large parameters remains a major challenge in practical HE

An example of the importance of good noise management

The graph-specific optimisation *lazy rescaling* reduces runtime by 8x.

- **Lazy rescaling**
 - Only after a **fully connected** or **convolution** layer
 - Skip if **no multiplications** left before decryption

Approximate Homomorphic Encryption over the Conjugate-invariant Ring Duhyeong Kim and Yongsoo Song

Table 5: Impact of lazy rescaling on CryptoNets runtime using $N = 2^{13}$, $L = 6$, with accuracy 98.95%.

Thread Count	Lazy Rescaling	Runtime	
		Amortized (ms)	Total (s)
1	✗	59.21	242.51 ± 3.69
1	✓	7.23	29.62 ± 0.63
24	✓	0.50	2.05 ± 0.11

Worst-case analysis

- Recall the canonical embedding $\sigma(a) = (a(\zeta_i))_{i \in (\mathbb{Z}/N\mathbb{Z})^*}$
- Use the **Canonical embedding norm**: $\|a\|^{\text{can}} = \|\sigma(a)\|_\infty$
- Variance of $\sigma(a)$ is nV_a

We use

$$\|a\|^{\text{can}} \leq 6\sqrt{n}\sqrt{V_a}, \quad (1)$$

and the following facts:

$$V_{a+b} = V_a + V_b$$

$$V_{\gamma a} = \gamma^2 V_a$$

$$V_{ab} = nV_a V_b.$$

BGV results: heuristic bounds vs. observed HElib noise growth

Enc			Add			Mult			ModSw		
P	I	\bar{x}	P	I	\bar{x}	P	I	\bar{x}	P	I	\bar{x}
34.0	35.0	41.1	33.0	34.0	40.2	14.0	17.0	26.0	-	-	-
88.0	89.0	97.9	87.0	88.0	97.0	67.0	70.0	82.4	38.0	39.0	38.1
196	197	209	195	196	209	174	177	194	146	147	150
415	416	433	414	415	432	392	395	416	365	366	373

Table 1: The observed mean \bar{x} of the noise budget in HElib ciphertexts in 10000 trials, with estimates of the noise budget obtained using Iliashenko heuristic analysis I and previous heuristic analysis P . Each row corresponds to a parameter set with $n \in \{2048, 4096, 8192, 16384\}$.

Average-case analysis for CKKS

- In CKKS, there are two sources of precision loss
 - Encoding noise
 - Encryption noise that never gets removed
- There is a way to detangle the two and provide an average-case analysis
- Track the noise variance all the way through a circuit via the Central Limit Theorem (CLT)
- Bound it at the end

Noise in the ring

$\log(N)$	$\log(q)$	Average	CLT
Addition noise.			
13	109	10.88	11.40
14	219	11.44	11.93
15	443	12.00	12.45
Multiplication noise.			
13	109	17.31	18.69
14	219	18.38	19.72
15	443	19.43	20.75

Table 2: Average bits of noise observed in the ring over 1000 trials in HEAAN, for $\alpha = 0.0001$ and $\Delta = 2^{40}$.

Results in the complex space

$\log(N)$	$\log(q)$	Average	CLT
Addition, complex error.			
13	109	-21.92	-22.55
14	219	-20.72	-21.52
15	443	-19.70	-20.49
Multiplication, complex error.			
13	109	-23.17	-21.51
14	219	-21.68	-19.92
15	443	-20.13	-18.72

Table 3: Average bits of error observed in the message space over 1000 trials in HEAAN, for $\alpha = 0.0001$ and $\Delta = 2^{40}$.

Evaluating functions homomorphically

What functions can we evaluate homomorphically?

- CKKS/ BGV/ BFV schemes are naturally limited to multiplication and addition

CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy; Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, John Wernsing

Chet: an optimizing compiler for fully-homomorphic neural- network inferencing; Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K., Maleki, S., Musuvathi, M., Mytkowicz, T

What functions can we evaluate homomorphically?

- CKKS/ BGV/ BFV schemes are naturally limited to multiplication and addition
- This can present challenges in some applications, such as Machine Learning

CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy; Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, John Wernsing

Chet: an optimizing compiler for fully-homomorphic neural- network inferencing; Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K., Maleki, S., Musuvathi, M., Mytkowicz, T

What functions can we evaluate homomorphically?

- CKKS/ BGV/ BFV schemes are naturally limited to multiplication and addition
- This can present challenges in some applications, such as Machine Learning
 - The activation function used in Neural Networks is typically non-linear

CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy; Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, John Wernsing

Chet: an optimizing compiler for fully-homomorphic neural- network inferencing; Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K., Maleki, S., Musuvathi, M., Mytkowicz, T

What functions can we evaluate homomorphically?

- CKKS/ BGV/ BFV schemes are naturally limited to multiplication and addition
- This can present challenges in some applications, such as Machine Learning
 - The activation function used in Neural Networks is typically non-linear
 - Examples include ReLu, $f(x) = \max(0, x)$

CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy; Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, John Wernsing

Chet: an optimizing compiler for fully-homomorphic neural- network inferencing; Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K., Maleki, S., Musuvathi, M., Mytkowicz, T

What functions can we evaluate homomorphically?

- CKKS/ BGV/ BFV schemes are naturally limited to multiplication and addition
- This can present challenges in some applications, such as Machine Learning
 - The activation function used in Neural Networks is typically non-linear
 - Examples include ReLu, $f(x) = \max(0, x)$
- Typically, we would approximate the function to evaluate

CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy; Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, John Wernsing

Chet: an optimizing compiler for fully-homomorphic neural- network inferencing; Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K., Maleki, S., Musuvathi, M., Mytkowicz, T

What functions can we evaluate homomorphically?

- CKKS/ BGV/ BFV schemes are naturally limited to multiplication and addition
- This can present challenges in some applications, such as Machine Learning
 - The activation function used in Neural Networks is typically non-linear
 - Examples include ReLu, $f(x) = \max(0, x)$
- Typically, we would approximate the function to evaluate
 - Usually via a Taylor Series approximation

CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy; Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, John Wernsing

Chet: an optimizing compiler for fully-homomorphic neural- network inferencing; Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K., Maleki, S., Musuvathi, M., Mytkowicz, T

What functions can we evaluate homomorphically?

- CKKS/ BGV/ BFV schemes are naturally limited to multiplication and addition
- This can present challenges in some applications, such as Machine Learning
 - The activation function used in Neural Networks is typically non-linear
 - Examples include ReLu, $f(x) = \max(0, x)$
- Typically, we would approximate the function to evaluate
 - Usually via a Taylor Series approximation
 - **CryptoNets uses $f(x) = x^2$, achieves 99% accuracy**

CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy; Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, John Wernsing

Chet: an optimizing compiler for fully-homomorphic neural- network inferencing; Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K., Maleki, S., Musuvathi, M., Mytkowicz, T

What functions can we evaluate homomorphically?

- CKKS/ BGV/ BFV schemes are naturally limited to multiplication and addition
- This can present challenges in some applications, such as Machine Learning
 - The activation function used in Neural Networks is typically non-linear
 - Examples include ReLu, $f(x) = \max(0, x)$
- Typically, we would approximate the function to evaluate
 - Usually via a Taylor Series approximation
 - CryptoNets uses $f(x) = x^2$, achieves 99% accuracy
 - CHET uses $f(x) = ax + b$, for $a, b \in \mathbb{R}$, achieves 81.5% accuracy (down from 84%)

CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy; Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, John Wernsing

Chet: an optimizing compiler for fully-homomorphic neural- network inferencing; Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K., Maleki, S., Musuvathi, M., Mytkowicz, T

Machine Learning Applications

Because of the limitations specified above, there are roughly two approaches:

- Modify the network to make it HE-friendly, and re-train

Side note: we mainly consider Neural Networks (NN) in this conversation, but other scenarios are more achievable for FHE [1].

[1]: Privacy-preserving semi-parallel logistic regression training with Fully Homomorphic Encryption; Sergiu Carpov, Nicolas Gama, Mariya Georgieva, and Juan Ramon Troncoso-Pastoriza

Machine Learning Applications

Because of the limitations specified above, there are roughly two approaches:

- Modify the network to make it HE-friendly, and re-train
 - Could lead to much better accuracy and overall performance as the network is now optimised for HE

Side note: we mainly consider Neural Networks (NN) in this conversation, but other scenarios are more achievable for FHE [1].

[1]: Privacy-preserving semi-parallel logistic regression training with Fully Homomorphic Encryption; Sergiu Carpov, Nicolas Gama, Mariya Georgieva, and Juan Ramon Troncoso-Pastoriza

Machine Learning Applications

Because of the limitations specified above, there are roughly two approaches:

- Modify the network to make it HE-friendly, and re-train
 - Could lead to much better accuracy and overall performance as the network is now optimised for HE
- **Modify the scheme or protocol**

Side note: we mainly consider Neural Networks (NN) in this conversation, but other scenarios are more achievable for FHE [1].

[1]: Privacy-preserving semi-parallel logistic regression training with Fully Homomorphic Encryption; Sergiu Carpov, Nicolas Gama, Mariya Georgieva, and Juan Ramon Troncoso-Pastoriza

Machine Learning Applications

Because of the limitations specified above, there are roughly two approaches:

- Modify the network to make it HE-friendly, and re-train
 - Could lead to much better accuracy and overall performance as the network is now optimised for HE
- Modify the scheme or protocol
 - No need to re-train your network, which is potentially a huge overhead

Side note: we mainly consider Neural Networks (NN) in this conversation, but other scenarios are more achievable for FHE [1].

[1]: Privacy-preserving semi-parallel logistic regression training with Fully Homomorphic Encryption; Sergiu Carpov, Nicolas Gama, Mariya Georgieva, and Juan Ramon Troncoso-Pastoriza

Machine Learning Applications

Because of the limitations specified above, there are roughly two approaches:

- Modify the network to make it HE-friendly, and re-train
 - Could lead to much better accuracy and overall performance as the network is now optimised for HE
- Modify the scheme or protocol
 - No need to re-train your network, which is potentially a huge overhead
- Some scenarios are easier than others (encrypted data vs encrypted model)

Side note: we mainly consider Neural Networks (NN) in this conversation, but other scenarios are more achievable for FHE [1].

[1]: Privacy-preserving semi-parallel logistic regression training with Fully Homomorphic Encryption; Sergiu Carpov, Nicolas Gama, Mariya Georgieva, and Juan Ramon Troncoso-Pastoriza

Machine Learning Applications

Because of the limitations specified above, there are roughly two approaches:

- Modify the network to make it HE-friendly, and re-train
 - Could lead to much better accuracy and overall performance as the network is now optimised for HE
- Modify the scheme or protocol
 - No need to re-train your network, which is potentially a huge overhead
- Some scenarios are easier than others (encrypted data vs encrypted model)
- Training in general is an open problem (backpropagation)

Side note: we mainly consider Neural Networks (NN) in this conversation, but other scenarios are more achievable for FHE [1].

[1]: Privacy-preserving semi-parallel logistic regression training with Fully Homomorphic Encryption; Sergiu Carpov, Nicolas Gama, Mariya Georgieva, and Juan Ramon Troncoso-Pastoriza

Interlude: Iterative Algorithms and Approximate Encryption

Iterative algorithms and approximate encryption

- In an exact scheme, we would only need to worry about correctness (and perhaps parameter selection)
- In the clear, we would exactly calculate the $(n + 1)^{\text{th}}$ estimate from the n^{th} to obtain a more accurate approximation
- With approximate encryption, the additional noise incurred by the $(n + 1)^{\text{th}}$ calculation may mean that calculating the $(n + 1)^{\text{th}}$ estimate doesn't provide an advantage over the n^{th} estimate
 - In fact, it may mean that your precision starts worsening instead of improving

Iterative algorithms and approximate encryption

For a given set of parameters, function, and data set, we want to identify the point at which the evaluation noise begins to interfere with the accurate bits of the calculated solution.

- Let $A^{(r)} = \|f(\mathbf{z}) - \mathbf{z}^{(r)}\|$ be the absolute error associated with iteration r in the clear,
- $\beta^{(r)}$ be a bound on the real noise in $\mathbb{C}^{N/2}$ at iteration r
- Want to identify the point at which


$$\log \beta^{(r)} > \log A^{(r)}$$

- If this critical point is correctly identified, the maximum possible accuracy that can be guaranteed is

$$-\log A^{(r-1)},$$

as by iteration r encryption noise has begun to interfere with the accuracy $-\log A^{(r)}$

Newton-Raphson



(c) $\log N = 15, \log \Delta = 35, a = 1, b = 15.$

Real-life implementation

- We have now reached the point where we will “soon” start deploying FHE
- This means we need to start thinking about
 - Lack of CCA security

- We have now reached the point where we will “soon” start deploying FHE
- This means we need to start thinking about
 - Lack of CCA security
 - By definition, ciphertexts are malleable

- We have now reached the point where we will “soon” start deploying FHE
- This means we need to start thinking about
 - Lack of CCA security
 - By definition, ciphertexts are malleable
 - Trusted Execution Environments (TEE)?

- We have now reached the point where we will “soon” start deploying FHE
- This means we need to start thinking about
 - Lack of CCA security
 - By definition, ciphertexts are malleable
 - Trusted Execution Environments (TEE)?
 - Key management?

- We have now reached the point where we will “soon” start deploying FHE
- This means we need to start thinking about
 - Lack of CCA security
 - By definition, ciphertexts are malleable
 - Trusted Execution Environments (TEE)?
 - Key management?
 - With eg AES you might use Hardware Security Modules (HSM); not practical for FHE


- We have now reached the point where we will “soon” start deploying FHE
- This means we need to start thinking about
 - Lack of CCA security
 - By definition, ciphertexts are malleable
 - Trusted Execution Environments (TEE)?
 - Key management?
 - With eg AES you might use Hardware Security Modules (HSM); not practical for FHE
 - ...?

HE as a Primitive

One can also think of HE as a primitive, which enables many constructions.

- Multi-Party Computation (MPC)
- Private Set Intersection (PSI)
- Privacy-Preserving Smart Contracts
- Private Set Union
- ...


PSI from FHE



Choose a random nonzero r .

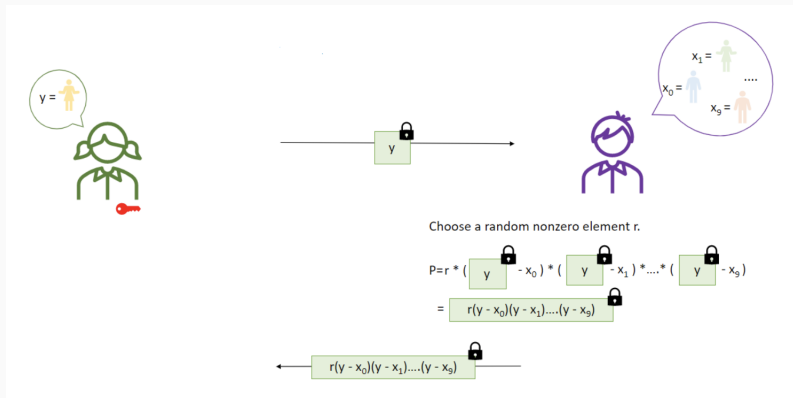
$$r * (y - x) = r(y - x)$$

The equation is displayed with a green box around "y" and a black padlock icon above it. The minus sign is followed by "x". The right side of the equation has a green box around "r(y - x)" and a black padlock icon above it.



<https://bit-ml.github.io/blog/post/private-set-intersection-an-implementation-in-python/>

PSI from FHE



<https://bit-ml.github.io/blog/post/private-set-intersection-an-implementation-in-python/>

Conclusion

Dimension n	bit-size t	sparse-subset-size s	big-set size S	big-set ratio R
512	380	15	512	2^{26}
2048	380	15	512	2^{102}
8192	380	15	547	2^{381}
32768	380	15	2185	2^{381}

Dimension n	bit-size t	# of ctexts in PK ($s \cdot c$)	PK size $\approx s \cdot c \cdot d $	keyGen	Recrypt
512	380	690	17 MByte	2.5 sec	6 sec
2048	380	690	69 MByte	41 sec	32 sec
8192	380	705	284 MByte	8.4 min	2.8 min
32768	380	1410	2.25 GByte	2.2 hour	31 min

Table 3: Parameters of the fully homomorphic scheme, as used for the public challenges.

Implementing Gentry's Fully Homomorphic Encryption Scheme, Craig Gentry and Shai Halevi <https://eprint.iacr.org/2010/520.pdf>.

Method	Acc. (%)	Lat. (s)	Thput. (im/s)	Protocol	Comm. (MB/im)
LoLa [10]	98.95	2.2	0.5	HE	
FHE-DiNN100 [8]	96.35	1.65	0.6	HE	
CryptoNets [24]	98.95	250	16.4	HE	
Faster CryptoNets [16]	98.7	39.1	210	HE	
nGraph-HE [6]	98.95	16.7	245	HE	
CryptoNets 3.2 [10]	98.95	25.6	320	HE	
nGraph-HE2	98.95	2.05	1,998	HE	
Chameleon [36]	99	2.24	1.0	HE-MPC	5.1
MiniONN [31]	98.95	1.28	2.4	HE-MPC	44
Gazelle [27]	98.95	0.03	33.3	HE-MPC	0.5
nGraph-HE2-ReLU	98.62	0.69	2,959	HE-MPC	0.03

nGraph-HE2: A High-Throughput Framework for Neural Network Inference on Encrypted Data, Fabian Boemer, Rosario Cammarota, **Anamaria Costache** and Casimir Wierzynski <https://eprint.iacr.org/2019/947.pdf>

Conclusion

- Some spectacular advances have been made since the first scheme in 2009
- In particular, this means that we are realistically looking at deploying FHE in the industry in the next “few years”
- There are some very exciting open problems, both theoretical and practical
- ... join us :)

Thank you!