

LOW COMMUNICATION THRESHOLD FULLY HOMOMORPHIC ENCRYPTION

ALAIN PASSELÈGUE & DAMIEN STEHLÉ

KOLKATA --- DECEMBER 13, 2024

Eprint 2024/1984



MAIN RESULTS

Contribution #1

Cryptanalysis of the BS23* Threshold-FHE with **moderate** decryption modulus

⇒ All known (general-purpose) Threshold-FHE's need **exponential** decr. modulus

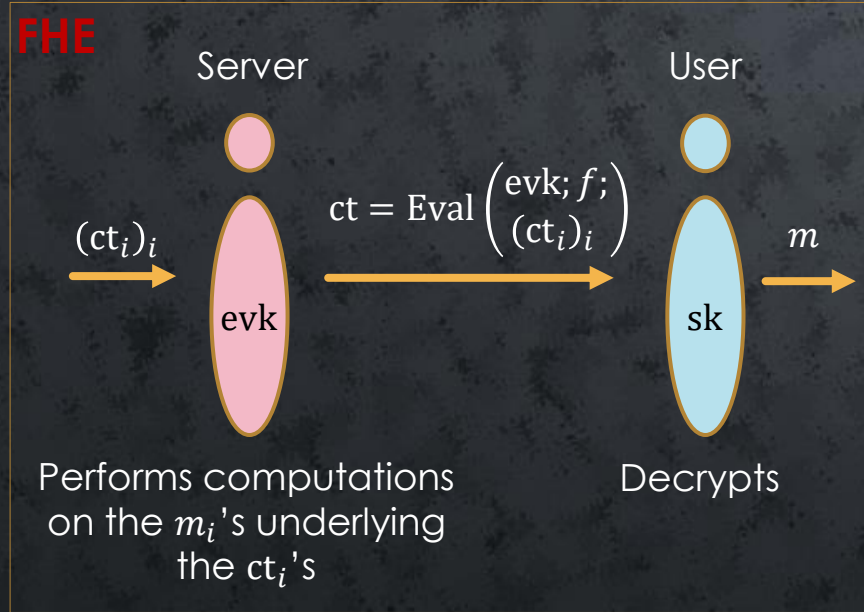
Contribution #2

Construction of a Threshold-FHE with **tiny** decryption modulus...

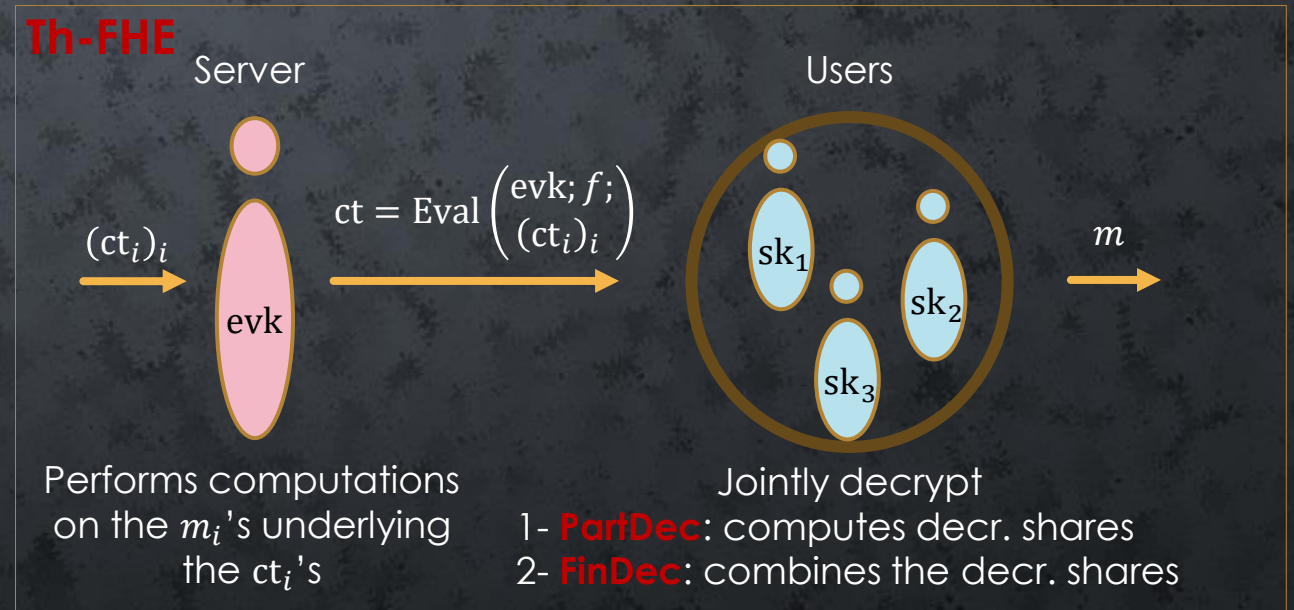
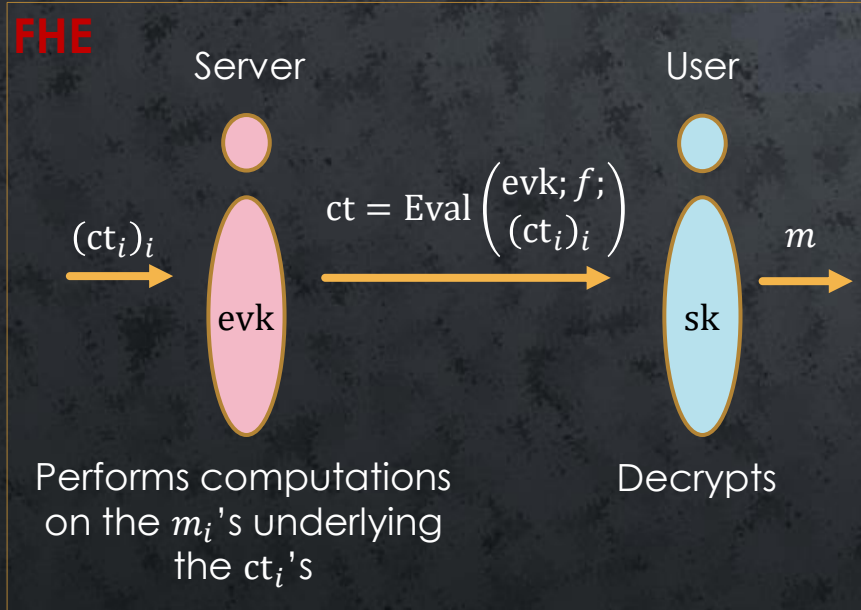
... for the (specific) case where the computing party is not corrupted.

Disclaimer: we only look at the N -out-of- N case

THRESHOLD-FHE



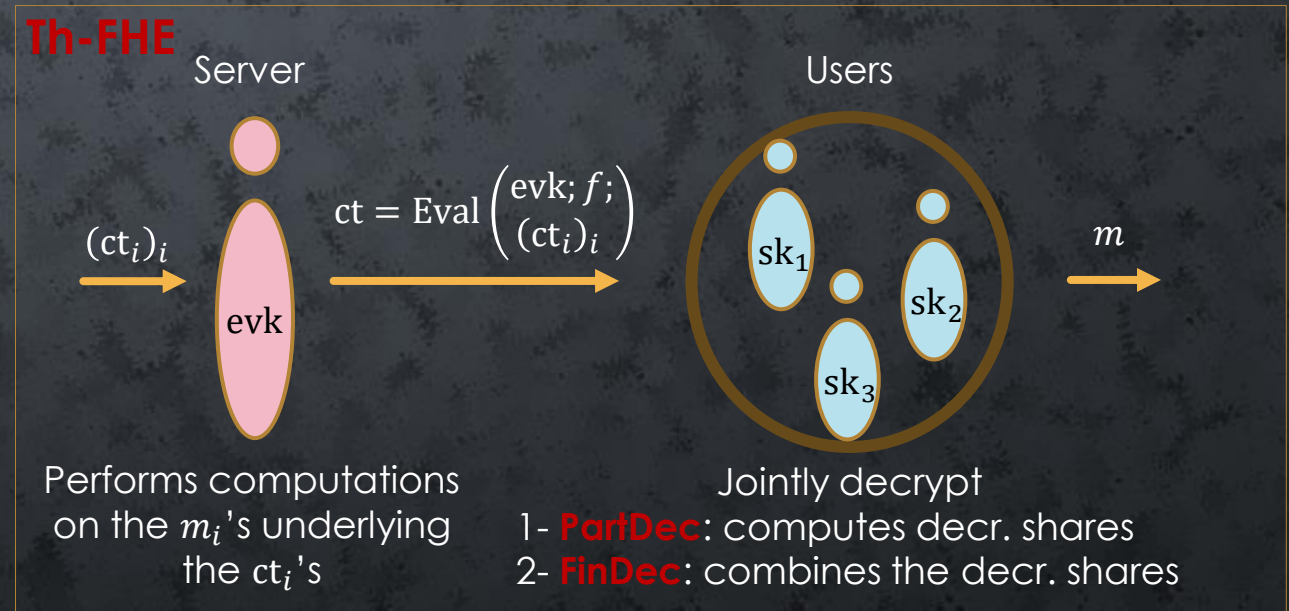
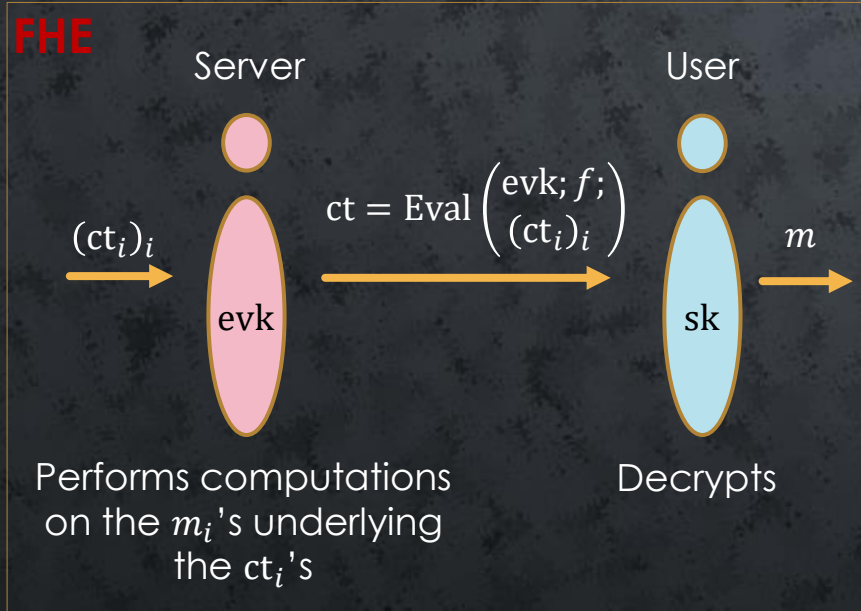
THRESHOLD-FHE



THRESHOLD-FHE

* G. Asharov, A. Jain, A. López-Alt, E. Tromer, V. Vaikuntanathan, D. Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. EUROCRYPT'12

** D. Boneh, R. Gennaro, S. Goldfeder, A. Jain, S. Kim, P. Rasmussen, A. Sahai. Threshold cryptosystems from threshold fully homomorphic encryption. CRYPTO'18



Th-FHE: sk is shared between users

- Protects sk
- Enables secure multi-party computations*
- Allows to thresholdize cryptographic constructions**

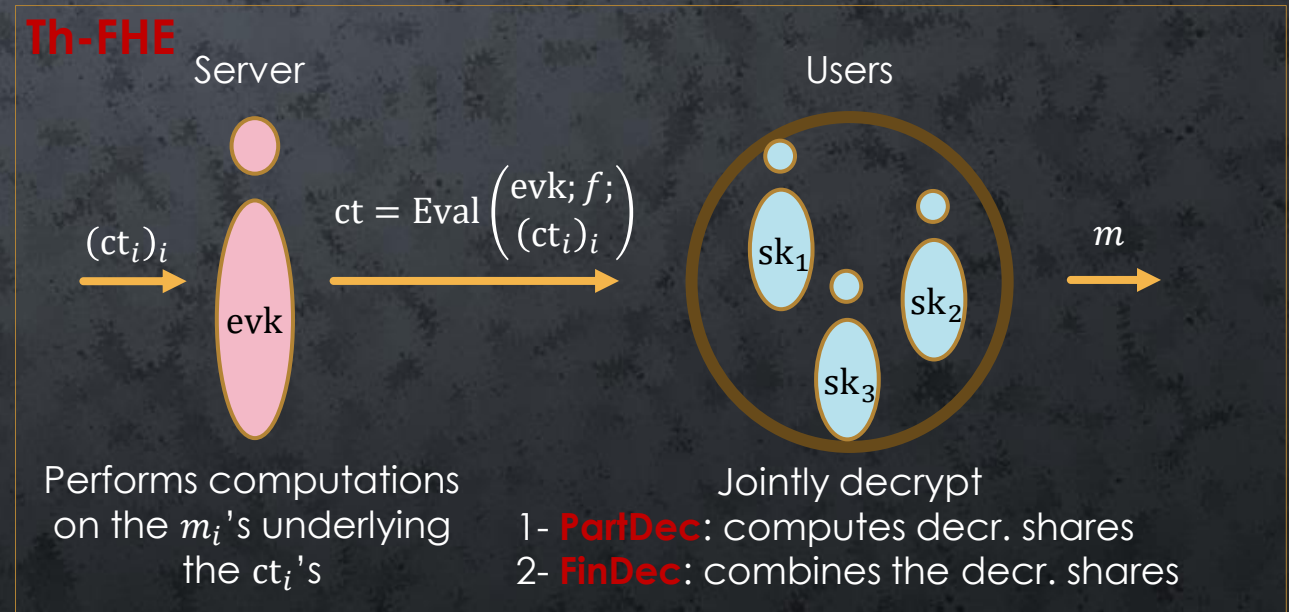
THRESHOLD-FHE: SECURITY (INFORMAL)

Adversary can:

- corrupt $N - 1$ users
- request encr. of ptxts
- request evaluations on generated ctxts
- request decr. of any generated ctxt (unless its ptxt trivially solves the challenge)

Adersary's challenge: distinguish between the encryptions of two ptxts of its choice

(we actually consider simulation-based security)



Side note: for $N = 1$ user, this matches the IND-CPA-D security notion from LM21*.

THRESHOLD-FHE: GENERAL DESIGN

Start from an FHE scheme, with ciphertexts of the form:

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e \ [q]$$

- a, b can be over \mathbb{Z}_q (LWE) or R_q (Ring-LWE), and e is small
- Ecd can be most/least/... significant bits

Split the key as

$$sk = \sum sk_i$$

PartDec (ct, sk_i)

$$sh_i := a \cdot sk_i + e_i$$

FinDec $(ct, (sh_i)_i)$

$$\text{Dcd} (\sum sh_i + b)$$

HOW LARGE SHOULD THE FLOODING BE?

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e [q]$$

Split the key as

$$sk = \sum sk_i$$

PartDec (ct, sk_i)

$$sh_i := a \cdot sk_i + e_i$$

FinDec $(ct, (sh_i)_i)$

$$\text{Dcd}(\sum sh_i + b)$$

The size of e_i drives the choice of q during decryption
⇒ drives amount of communication in decryption
(if need be, we can switch to a large q just before decryption*)

HOW LARGE SHOULD THE FLOODING BE?

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e \ [q]$$

Split the key as

$$sk = \sum sk_i$$

PartDec (ct, sk_i)

$$sh_i := a \cdot sk_i + e_i$$

FinDec (ct, $(sh_i)_i$)

$$\text{Dcd}(\sum sh_i + b)$$

The size of e_i drives the choice of q during decryption
⇒ drives amount of communication in decryption
(if need be, we can switch to a large q just before decryption*)

No e_i : $e_i = 0$

Adversary can recover sk_i

Exponential e_i : $|e_i| \geq 2^\lambda \cdot |e|$

We can simulate the adversary's view

HOW LARGE SHOULD THE FLOODING BE?

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e \ [q]$$

Split the key as

$$sk = \sum sk_i$$

PartDec (ct, sk_i)

$$sh_i := a \cdot sk_i + e_i$$

FinDec $(ct, (sh_i)_i)$

$$\text{Dcd}(\sum sh_i + b)$$

No e_i : $e_i = 0$

Adversary can recover sk_i

Exponential e_i : $|e_i| \geq 2^\lambda \cdot |e|$

We can simulate the adversary's view

Very small e_i : $|e_i| \approx \text{poly}(\lambda) \cdot |e|$

Adversary can "average-out" e_i in sh_i

What about moderate e_i ?

$$|e_i| \approx \text{poly}(Q_{dec}) \cdot |e|$$

(Q_{dec} is the number of decr. queries)

CAN MODERATE FLOODING WORK?

* S. Agrawal, D. Stehlé, A. Yadav. Round-optimal lattice-based threshold signatures, revisited. ICALP'22

** B. Li, D. Micciancio, M. Schultz, J. Sorrell. Securing approximate homomorphic encryption using differential privacy. CRYPTO'22

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e \ [q]$$

Split the key as

$$sk = \sum sk_i$$

PartDec (ct, sk_i)

$$sh_i := a \cdot sk_i + e_i$$
$$|e_i| \approx \text{poly}(Q_{dec}) \cdot |e|$$

FinDec $(ct, (sh_i)_i)$

$$\text{Dcd}(\sum sh_i + b)$$

ASY22*: When used to thresholdize a signature scheme, this can be proved secure using Rényi Divergence

LMSS22** attack, e.g. using BFV:

- Encrypt $(10,10)$ or $(0,0)$
- Perform an inner product with $(1, -1)$
- In both cases, the result is 0
- But the noise is larger for $(10,10)$

⇒ Gives a $\text{poly}\left(\frac{1}{Q_{dec}}\right)$ dist. advantage

THE BS23 APPROACH

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e \ [q]$$

Split the key as

$$sk = \sum sk_i$$

PartDec (ct, sk_i)

$$sh_i := a \cdot sk_i + e_i \\ |e_i| \approx \text{poly}(Q_{dec}) \cdot |e|$$

FinDec $(ct, (sh_i)_i)$

$$\text{Dcd}(\sum sh_i + b)$$

ASY22: When used to thresholdize a signature scheme, this can be proved secure using Rényi Divergence

BS23* (informal): Assuming the FHE scheme is **circuit-private**, then the threshold FHE scheme is secure with moderate noise

LMSS22: The decryption noise may carry **information on past computations**, including the challenge plaintexts

Circuit privacy: the distribution of the decryption noise does not depend on past computations, even if sk is given to the adversary

CONTRIBUTION #1: CRYPTANALYSIS OF BS23

We are given $ct^* = (a^*, b^*)$: $a^* \cdot sk + b^* = \text{Ecd}(m_\beta) + e \ [q]$

We want to distinguish $\beta = 0$ from $\beta = 1$

Assumption 1: the scheme allows “Rescale”

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e \ [q]$$



$$ct' = (a', b') = \left(\left\lfloor \frac{q'}{q} a \right\rfloor, \left\lfloor \frac{q'}{q} b \right\rfloor \right) \ [q'] \quad \text{with } q' \ll q$$

We have

$$a' \cdot sk + b' \approx \text{Ecd}(m) + e_{rnd} \cdot sk,$$

where $e_{rnd} = \left\{ \frac{q'}{q} a \right\}$ is **known**

CONTRIBUTION #1: CRYPTANALYSIS OF BS23

We are given $ct^* = (a^*, b^*)$: $a^* \cdot sk + b^* = \text{Ecd}(m_\beta) + e \ [q]$

We want to distinguish $\beta = 0$ from $\beta = 1$

Assumption 1: the scheme allows “Rescale”

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e \ [q]$$



$$ct' = (a', b') = \left(\left\lfloor \frac{q'}{q} a \right\rfloor, \left\lfloor \frac{q'}{q} b \right\rfloor \right) \ [q'] \quad \text{with } q' \ll q$$

We have

$$a' \cdot sk + b' \approx \text{Ecd}(m) + e_{rnd} \cdot sk,$$

where $e_{rnd} = \left\{ \frac{q'}{q} a \right\}$ is **known**

Assumption 2: “nice” homomorphic mult. noise

$$ct_1 = (a_1, b_1): a_1 \cdot sk + b_1 = \text{Ecd}(m_1) + e_1 \ [q]$$

$$ct_2 = (a_2, b_2): a_2 \cdot sk + b_2 = \text{Ecd}(m_2) + e_2 \ [q]$$



$$ct_x = (a_x, b_x): a_x \cdot sk + b_x = \text{Ecd}(m_1 \cdot m_2) + e_x \ [q']$$

with

$$e_x \approx m_1 \cdot e_2 + m_2 \cdot e_1$$

BFV and CKKS can be parametrized to fit

CONTRIBUTION #1: CRYPTANALYSIS OF BS23

We are given $ct^* = (a^*, b^*)$: $a^* \cdot sk + b^* = \text{Ecd}(m_\beta) + e \ [q]$

We want to distinguish $\beta = 0$ from $\beta = 1$

If Eval ends up with $\text{Enc}(A)$,
we post-process as follows:

$\text{Enc}(1)$

↓ Rescale

$\text{Enc}(1)$

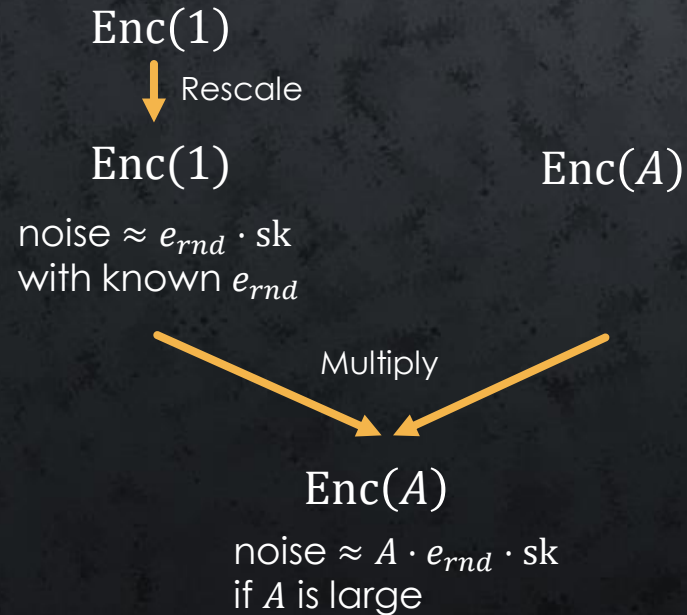
noise $\approx e_{rnd} \cdot sk$
with known e_{rnd}

CONTRIBUTION #1: CRYPTANALYSIS OF BS23

We are given $ct^* = (a^*, b^*)$: $a^* \cdot sk + b^* = \text{Ecd}(m_\beta) + e \ [q]$

We want to distinguish $\beta = 0$ from $\beta = 1$

If Eval ends up with $\text{Enc}(A)$,
we post-process as follows:

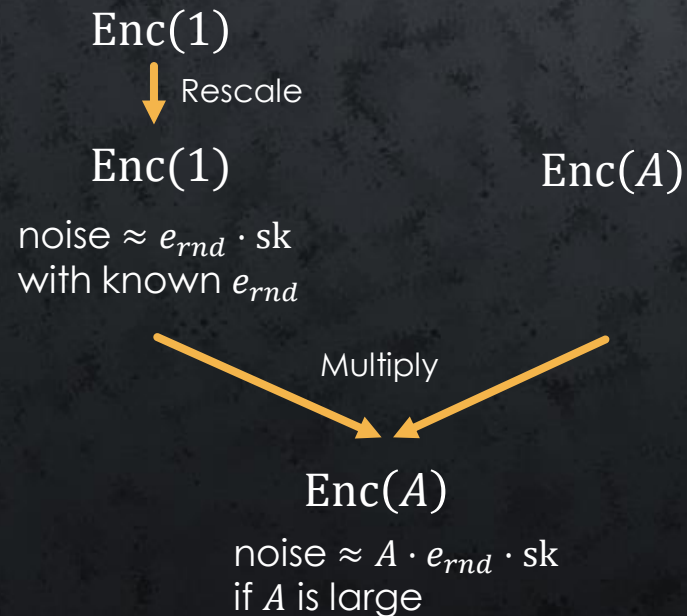


CONTRIBUTION #1: CRYPTANALYSIS OF BS23

We are given $ct^* = (a^*, b^*)$: $a^* \cdot sk + b^* = \text{Ecd}(m_\beta) + e \ [q]$

We want to distinguish $\beta = 0$ from $\beta = 1$

If Eval ends up with $\text{Enc}(A)$,
we post-process as follows:



1. If the initial scheme is circuit-private, then so is the modified scheme
2. Request encryption and decryption of $\text{MSB}(a^*)$
3. Recover $e_x \approx \text{MSB}(a^*) \cdot e_{rnd} \cdot sk$
4. Compute $e_x + e_{rnd} \cdot b^* \approx e_{rnd} \cdot \text{Ecd}(m_\beta)$
5. As the decr. modulus is small, we can distinguish

FORGETTING HISTORY REQUIRES RANDOMNESS

Deeper issue with the BS23 approach:

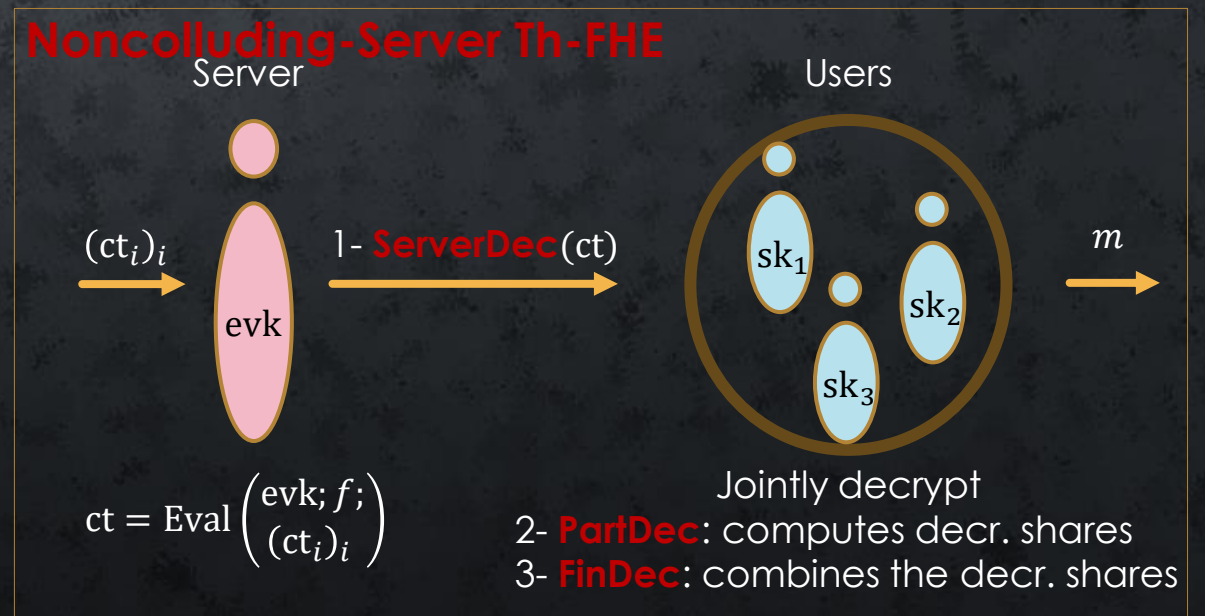
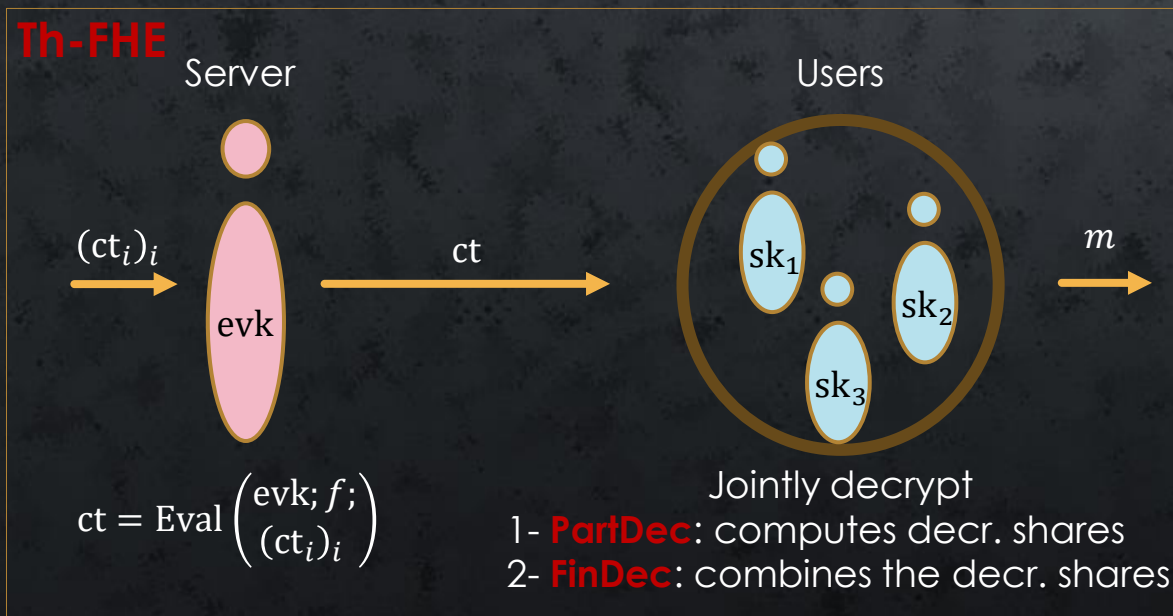
- Current circuit-privacy techniques require the server to **inject randomness**
- But the server is potentially a corrupted user \Rightarrow **not random to the adversary**

FORGETTING HISTORY REQUIRES RANDOMNESS

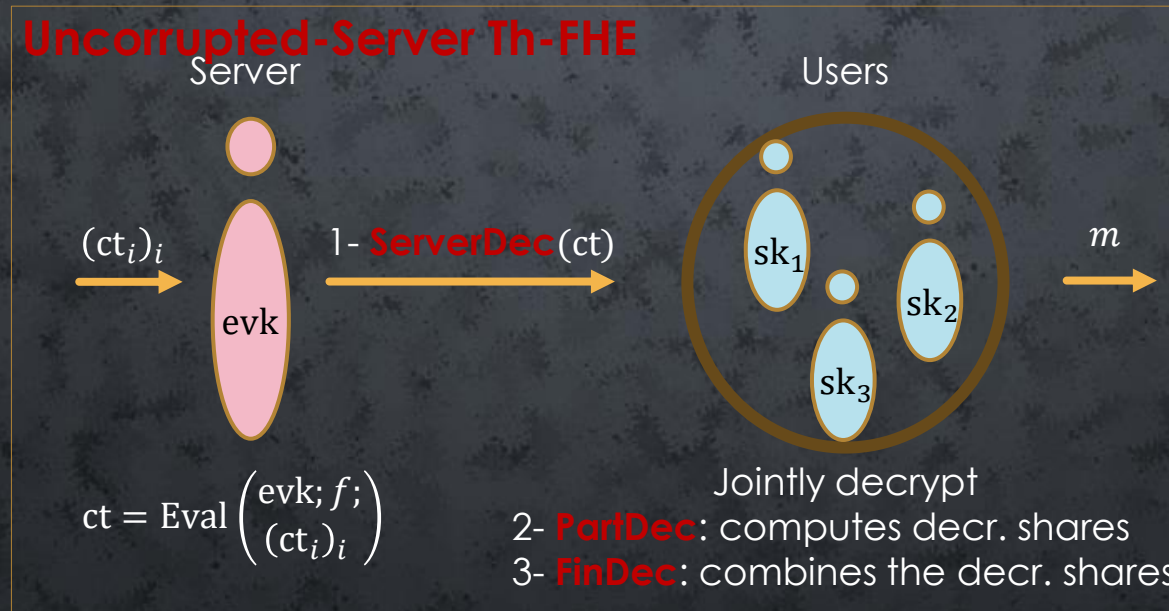
Deeper issue with the BS23 approach:

- Current circuit-privacy techniques require the server to **inject randomness**
- But the server is potentially a corrupted user \Rightarrow **not random to the adversary**

We propose an **noncolluding-server variant of Threshold-FHE**



DOES THE MODEL MAKE SENSE?



This depends on applications!

- OK if the group of users externalizes the computation to an outsider server
⇒ Somewhat trusted third party (may eavesdrop, may not collude)
- Not OK for the universal thresholdizer, which requires Eval to be deterministic

CONTRIBUTION #2: DOUBLE FLOOD & ROUND

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e \ [q]$$

Split the key as

$$sk = \sum sk_i$$

ServerDec (ct)

Add Enc(0) to ct
Add exponential flooding to b -part
Rdm-Rescale to a **poly(λ) modulus**

PartDec (ct, sk_i)

$$sh_i := a \cdot sk_i + e_i$$

Using **tiny** e_i

FinDec (ct, $(sh_i)_i$)

$$\text{Dcd}(\sum sh_i + b)$$

CONTRIBUTION #2: DOUBLE FLOOD & ROUND

$$ct = (a, b): a \cdot sk + b = \text{Ecd}(m) + e \ [q]$$

Split the key as

$$sk = \sum sk_i$$

ServerDec (ct)

Add Enc(0) to ct
Add exponential flooding to b -part
Rdm-Rescale to a **poly(λ) modulus**

PartDec (ct, sk_i)

$$sh_i := a \cdot sk_i + e_i$$

Using **tiny** e_i

FinDec (ct, $(sh_i)_i$)

$$\text{Dcd}(\sum sh_i + b)$$

- Everything that users get and send is with a **poly(λ) modulus**
- This requires exponential flooding, but only internally to the server
- Proof technique closely related to MS23*

WRAP-UP

Contribution #1

Cryptanalysis of the BS23* Threshold-FHE with **moderate** decryption modulus

⇒ All known (general-purpose) Threshold-FHE's need **exponential** decr. modulus

Contribution #2

Construction of a Threshold-FHE with **tiny** decryption modulus...

... for the (specific) case where the server is not colluding.

Open problems:

- Can we get general-purpose Threshold-FHE with $\text{poly}(Q_{dec})$ decryption modulus?
- Can we weaken the noncolluding-server assumption?

QUESTIONS?

Eprint 2024/1984

{ alain.passelegue, damien.stehle } @ cryptolab.co.kr