#### Transistor

a TFHE-friendly stream cipher

Jules Baudrin, Sonia Belaïd, Nicolas Bon, Christina Boura, Anne Canteaut, Gaëtan Leurent, Pascal Paillier, Léo Perrin, Matthieu Rivain, Yann Rotella, and Samuel Tap



WE INNOVATE TO SECURE YOUR BUSINESS

29th October 2024

- FHE and the bandwith problem
- Transciphering
- TFHE and a wishlist
- Design of Transistor

Client









Client











Client











Client













Client











Client





Client Server 0 Enc<sub>FHE</sub> A  $\mathsf{Dec}_\mathsf{FHE}$ 1 0

Example: Layout of a TFHE ciphertext.

Example: Layout of a TFHE ciphertext. To encrypt one bit m:

$$c = (a_0, \ldots, a_{n-1}, b)$$

Example: Layout of a TFHE ciphertext. To encrypt one bit m:

$$c = (a_0, \ldots, a_{n-1}, b)$$

(n+1) elements of  $\mathbb{Z}_q$ , with  $q = 2^{64}$  and  $500 \le n \le 1500$ . Roughly, expansion factor of  $1000 \times 64 = 64000$ .

Example: Layout of a TFHE ciphertext. To encrypt one bit m:

$$c = \left(\underbrace{a_0, \dots, a_{n-1}}_{\mathsf{Random}}, b\right)$$

(n+1) elements of  $\mathbb{Z}_q$ , with  $q = 2^{64}$  and  $500 \le n \le 1500$ . Roughly, expansion factor of  $1000 \times 64 = 64000$ .

Possible to generate the  $a_i$ 's with a seed: expansion factor comes down to 128 + 64 = 192.

Example: Layout of a TFHE ciphertext. To encrypt one bit m:

$$c = \left(\underbrace{a_0, \dots, a_{n-1}}_{\mathsf{Random}}, b\right)$$

(n + 1) elements of  $\mathbb{Z}_q$ , with  $q = 2^{64}$  and  $500 \le n \le 1500$ . Roughly, expansion factor of  $1000 \times 64 = 64000$ .

Possible to generate the  $a_i$ 's with a seed: expansion factor comes down to 128 + 64 = 192.

#### How to reduce this bandwith overhead ?

- FHE and the bandwith problem
- Transciphering
- TFHE and a wishlist
- Design of Transistor

Client



















Client  $\mathsf{Enc}_{\mathsf{Sym}}$ Enc<sub>FHE</sub>



# Which symmetric cipher ?

Using standard ciphers ?

- Evaluation of AES is too slow due to large Sbox size
- Lightweight ciphers such as ASCON, PRESENT, ... more promising, but still not fast enough

# Which symmetric cipher ?

Using standard ciphers ?

- Evaluation of AES is too slow due to large Sbox size
- Lightweight ciphers such as ASCON, PRESENT, ... more promising, but still not fast enough

#### Or FHE-tailored ones ?

- Active line of research
- But often relies on innovative primitives (i.e. often broken)

Using standard ciphers ?

- Evaluation of AES is too slow due to large Sbox size
- Lightweight ciphers such as ASCON, PRESENT, ... more promising, but still not fast enough

Or FHE-tailored ones ?

- Active line of research
- But often relies on innovative primitives (i.e. often broken)

With Transistor, we look for the best of both worlds (fast in FHE and secure)

- FHE and the bandwith problem
- Transciphering
- TFHE and a wishlist
- Design of Transistor

• Plaintext space :  $\mathbb{Z}_p$  with p of a few bits.

- Plaintext space :  $\mathbb{Z}_p$  with p of a few bits.
- Secret Key:  $\vec{s} = (s_0, \dots, s_{n-1}) \in \{0, 1\}^n$

- Plaintext space :  $\mathbb{Z}_p$  with p of a few bits.
- Secret Key:  $\vec{s} = (s_0, \dots, s_{n-1}) \in \{0, 1\}^n$
- Encryption:  $\vec{c} = (a_0, \ldots, a_{n-1}, b = \sum_{i=0}^{n-1} a_i \cdot s_i + \frac{q}{p} \cdot m + e) \in \mathbb{Z}_q^{n+1}$  with  $\vec{a}$  random and e a small Gaussian noise.

- Plaintext space :  $\mathbb{Z}_p$  with p of a few bits.
- Secret Key:  $\vec{s} = (s_0, \dots, s_{n-1}) \in \{0, 1\}^n$
- Encryption:  $\vec{c} = (a_0, \ldots, a_{n-1}, b = \sum_{i=0}^{n-1} a_i \cdot s_i + \frac{q}{p} \cdot m + e) \in \mathbb{Z}_q^{n+1}$  with  $\vec{a}$  random and e a small Gaussian noise.
- Linear homomorphism:
  - Sum of ciphertexts
  - Multiplication by a plaintext

Very fast. Increase the noise.

- Plaintext space :  $\mathbb{Z}_p$  with p of a few bits.
- Secret Key:  $\vec{s} = (s_0, \dots, s_{n-1}) \in \{0, 1\}^n$
- Encryption:  $\vec{c} = (a_0, \dots, a_{n-1}, b = \sum_{i=0}^{n-1} a_i \cdot s_i + \frac{q}{p} \cdot m + e) \in \mathbb{Z}_q^{n+1}$  with  $\vec{a}$  random and e a small Gaussian noise.
- Linear homomorphism:
  - Sum of ciphertexts
  - Multiplication by a plaintext
  - Very fast. Increase the noise.
- Programmable Bootstrapping (PBS):
  - Evaluates a Look-Up Table
  - **Resets** the noise to a nominal level

Very slow. Gets even slower as p increases.

- Plaintext space :  $\mathbb{Z}_p$  with p of a few bits.
- Secret Key:  $\vec{s} = (s_0, \dots, s_{n-1}) \in \{0, 1\}^n$
- Encryption:  $\vec{c} = (a_0, \dots, a_{n-1}, b = \sum_{i=0}^{n-1} a_i \cdot s_i + \frac{q}{p} \cdot m + e) \in \mathbb{Z}_q^{n+1}$  with  $\vec{a}$  random and e a small Gaussian noise.
- Linear homomorphism:
  - Sum of ciphertexts
  - Multiplication by a plaintext
  - Very fast. Increase the noise.
- Programmable Bootstrapping (PBS):
  - Evaluates a Look-Up Table
  - **Resets** the noise to a nominal level
  - Very slow. Gets even slower as p increases.

#### • Negacyclicity problem:

- If p is **even**, restricts some homomorphic operations.
- Disappears when p is **odd**

# Our wishlist for a TFHE-friendly cipher

# Our wishlist for a TFHE-friendly cipher

• A prime field

- A prime field
- As little non-linear operations as possible

- A prime field
- As little non-linear operations as possible
- Keep noise growth acceptable

- FHE and the bandwith problem
- Transciphering
- TFHE and a wishlist
- Design of Transistor

# Design of Transistor

Prime field:  $\mathbb{F}_{17}$ 



(a) General structure (rectangles correspond to registers).



Figure: A high level view of Transistor.

#### Our wishlist

- A prime field
- As little non-linear operations as possible
- Smallest noise growth possible

# Design of Transistor

Prime field:  $\mathbb{F}_{17}$ 



(a) General structure (rectangles correspond to registers).



Figure: A high level view of Transistor

#### Our wishlist

- A prime field
- As little non-linear operations as possible
- Smallest noise growth possible

# Design of Transistor

Prime field:  $\mathbb{F}_{17}$ 



(a) General structure (rectangles correspond to registers).



Figure: A high level view of Transistor.

#### Our wishlist

- A prime field
- As little non-linear operations as possible
- Smallest noise growth possible

# Design of Transistor

Prime field:  $\mathbb{F}_{17}$ 



(a) General structure (rectangles correspond to registers).



Figure: A high level view of Transistor.

We opted for a  $4 \times 4$  Maximum Distance Separable (MDS) to ensure optimal diffusion. The matrix we chose is

$$M = \begin{bmatrix} 2 & 1 & 1 & 1 \\ 1 & -1 & 1 & -2 \\ 1 & 1 & -2 & -1 \\ 1 & -2 & -1 & 1 \end{bmatrix} .$$
(1)

We verified that there is no MDS matrix in  $\mathbb{F}_{17}$  with coefficients in  $\{-1, 1\}$  by exhaustively testing all such matrices. By testing all matrices with coefficients in  $\{-2, -1, 1, 2\}$ , we found a total of 30 720 MDS matrices with an  $\ell_2$ -norm of 7. We selected M for its symmetries, particularly because it is its own transpose.

# Design of Transistor

Prime field:  $\mathbb{F}_{17}$ 



Figure: A high level view of Transistor.

# Design of Transistor





• Naive approach would be to maintain an encrypted state, and update it by computing a linear combination with the feedback coefficients.



- Naive approach would be to maintain an encrypted state, and update it by computing a linear combination with the feedback coefficients.
- However, this method would cause the noise in the state to accumulate over time



- Naive approach would be to maintain an encrypted state, and update it by computing a linear combination with the feedback coefficients.
- However, this method would cause the noise in the state to accumulate over time
- Solution: Computing on the fly the coefficients of the linear combination in clear



- Naive approach would be to maintain an encrypted state, and update it by computing a linear combination with the feedback coefficients.
- However, this method would cause the noise in the state to accumulate over time
- Solution: Computing on the fly the coefficients of the linear combination in clear

# The noise variance in the output of the silent LFSR remains stable over time, without using any PBS.





#### The LFSR are loaded with fresh ciphertexts of noise $\sigma_{\text{fresh}}$



The output of silent LFSRs have noise

$$\sigma_{\mathcal{K}}^2 \leq |\mathcal{K}| \cdot \left(\frac{p-1}{2}\right)^2 \cdot \sigma_{\mathsf{fresh}}^2 \quad \text{and} \quad \sigma_{\mathcal{W}}^2 \leq |\mathcal{W}| \cdot \left(\frac{p-1}{2}\right)^2 \cdot \sigma_{\mathsf{fresh}}^2$$



The output of a PBS has noise **independent** from the noise of the input. We just have:

 $\sigma_{\rm PBS} \gg \sigma_{\rm fresh}$ 



ShiftRows simply reorders the ciphertexts, thus adds no noise.



The variance is multiplied by the norm of the diffusion matrix

$$\sigma_{\rm MC}^2 = L_{\rm MC}^2 \cdot \sigma_{\rm PBS}^2.$$



Variances are summed.



There are two "hot spots" where things can go wrong:

• Right before the PBS (risk of false bootstrapping result)



There are two "hot spots" where things can go wrong:

- Right before the PBS (risk of false bootstrapping result)
- At the output of the cipher (risk of producing a keystream of noise too high, thus unusable)



As  $\sigma_{MC} \gg \sigma_{\mathcal{K}}^2$  and  $\sigma_{PBS} \gg \sigma_{\mathcal{W}}^2$ , the noise produced by the LFSRs is negligible.

#### Takeaway 1: No Restriction on the size of the LFSRs



Takeaway 2 : Dimensioning the TFHE parameters for Transistor can be reduced to select parameters for a simple PBS

Table: Execution timings of FRAST and Transistor.

Cipher	$p_{err}$	Setup	Latency	Throughput
FRAST	$2^{-80}$	25 s (8 threads)	6.2 s	20.66 bits/s
Transistor	$2^{-128}$	No	251 ms	65.10 bits/s

A quick tour of the security arguments:

- Time-Memory-Data Trade-Offs put a limit on the length of the keystream we can generate with a single key:  $2^{31}$ .
- Guess-and-Determine: as one quarter of the state is used to generate the keystream, one has to guess  $\frac{3}{4}$  of the LFSR, which puts a lower bound on its size.  $|\mathcal{K}| = 64$  and  $\mathcal{W}| = 32$
- The paper demonstrates that **Three consecutive outputs are statistically independent of the secret key**. The proof only requires the matrix of MixColumns to be MDS instead of automatic search methods like in Rocca.
- Linear approximation, Fast Correlation attacks, Algebraic Attacks