At Last! A Homomorphic AES Evaluation in Less than 30 Seconds by Means of TFHE [*]

Daphné Trama, Pierre-Emmanuel Clet, Aymen Boudguiga, and Renaud Sirdey

Université Paris-Saclay, CEA-List, Palaiseau, France
daphne.trama@cea.fr
pierre-emmanuel.clet@cea.fr
aymen.boudguiga@cea.fr
renaud.sirdey@cea.fr

Abstract. Since the pioneering work of Gentry, Halevi, and Smart in 2012 [16], the state of the art on transciphering has moved away from work on AES to focus on new symmetric algorithms that are better suited for a homomorphic execution. Yet, with recent advances in homomorphic cryptosystems, the question arises as to where we stand today. Especially since AES execution is the application that may be chosen by NIST in the FHE part of its future call for threshold encryption. In this paper, we propose an AES implementation using TFHE programmable bootstrapping which runs in less than a minute on an average laptop. We detail the transformations carried out on the original AES code to lead to a more efficient homomorphic evaluation and we also give several execution times on different machines, depending on the type of execution (sequential or parallelized). These times vary from 4.5 minutes (resp. 54 secs) for sequential (resp. parallel) execution on a standard laptop down to 28 seconds for a parallelized execution over 16 threads on a multi-core workstation.

Keywords: AES · Fully Homomorphic Encryption · Transciphering · TFHE · Programmable Bootstrapping.

1 Introduction

With recent advances in FHE, is a homomorphic AES still as impractical as it was ten years ago? The work of Gentry, Halevi, and Smart [16] in 2012 pushed research towards new symmetric cryptosystems designed primarily to be faster to evaluate over FHE, paving the way for transciphering. Indeed, they performed an AES-128 homomorphic evaluation with BGV using HElib with now obsolete parameters that did not allow bootstrapping. They thus obtained an execution time of 4.1 minutes, but without allowing further operations on the final ciphertext. With bootstrapping, allowing calculation after the homomorphic execution of the AES, their runtime grew to 17.5 minutes. So neither of these two approaches could be used in practice. Several teams then decided to create new symmetric cryptosystems, whose encryption operations

were specifically chosen to be more rapidly executed in the homomorphic domain. As of today, there are many proposals, from block ciphers (LowMC [1], PRINCE [5], CHAGHRI [2]) to stream ciphers (Elisabeth [13], PASTA [15], Kreyvium [7]). Each comes with its pros and cons. For instance, PRINCE [5] is a block cipher especially created to be lightweight and, although it was initially proposed independently of Gentry's breakthrough on FHE, has a number of desirable properties with respect to homomorphic execution: a moderate number of rounds, small depth (for a block-cipher) and a low gate count/footprint of the decryption and encryption functions. PRINCE was one of the first symmetric algorithms for which an FHE execution attempt was done [19]. LowMC [1], on the other hand, is the first block cipher explicitly designed with FHE and MPC in mind. Although demonstrating competitive FHE execution performances at the time of proposal, its design was intrinsically bit-oriented while the FHE state of the art has moved away from bit-level FHE operations due to the relative inefficiency of this latter approach. In 2022 Ashur, Mahzoun, and Toprakhisar presented CHAGHRI [2], an FHE-friendly block cipher enabling efficient transciphering in BGV-like schemes. A complete CHAGHRI circuit can be implemented using 16 multiplications, 48 Frobenius automorphisms, and 32 rotations. The authors implemented it with HElib in order to compare it with Gentry et al., work on AES. Although their implementation is claimed to be 63% faster than [16], an attack on CHAGHRI has recently been proposed [20]. Also introduced in 2022, Elisabeth is a family of stream ciphers especially designed to be efficient for Hybrid Homomorphic Encryption (HHE). The authors use TFHE and propose a Rust implementation (using the Concrete library) of `Elisabeth-4`, that is to say, a cryptosystem in which inputs are on 4 bits. So it would take 32 executions of the cipher to obtain a 128-bits ciphertext. Before Elisabeth, the PASTA cryptosystem, implemented with BGV/BFV proposed an optimized cipher for integer HHE use cases. They also benchmark several HHE schemes, using the HElib library. But the use of a non-bootstrapping-able scheme limits the number of operations to be further performed on the ciphertexts. Kreyvium is a stream cipher, which is a variant of Trivium [6] (a streamcipher belonging to the eSTREAM portfolio). The main motivation for introducing Kreyvium was to present an FHE-friendly symmetric primitive with 128-bits of security. Additionally, the state of the art also includes homomorphic evaluation of several variants of the Grain-128 stream cipher by means of TFHE either in gate-bootstrapping mode or exploiting its functional bootstrapping capabilities [4, 3]. Often compared to the homomorphic execution times of AES as a guarantee of efficiency, none of these cryptosystems has been standardized (with the notable exception of Grain-128 which was a finalist in the recent NIST competition on lightweight cryptography). Yet, an "efficient-by-FHE-standards" homomorphic AES execution remains interesting for the research community working on transciphering, even if it does not bring any revolution. This is especially so, since AES execution may be the application chosen by NIST in the FHE part of its future call for proposals on threshold encryption[1].

Contribution– In this paper, we propose an AES implementation using TFHE programmable bootstrapping, which runs in less than a minute on a standard laptop PC. We first discuss the modifications carried out on the original AES code to lead

---

[1] `https://csrc.nist.gov/Projects/threshold-cryptography`.

to an efficient homomorphic evaluation. Then we give details about the benchmark made to determine which decomposition basis to use to have a faster evaluation of the algorithm. We finally provide experimental execution times on different machines, depending on the type of execution (sequential or parallelized).

Paper Ogranization– This paper is organized as follows: Section 2 reviews the basics of the TFHE cryptosystem and gives the necessary details of the tree-based method for bootstrapping with multi-input ciphertexts and its optimization with multi-value bootstrapping. Section 3 gives a brief reminder on the AES. Section 4 provides a detailed exposition of our approaches to transform the original AES code and implement the most optimized and efficient version of it with TFHE programmable bootstrapping. Section 5 presents the performances and results of our methods.

2   TFHE Preliminaries

2.1   Notations

Let $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ be the real torus, that is to say, the additive group of real numbers modulo 1 ($\mathbb{R} \mod 1$). We will denote by $\mathbb{T}_N[X]^n$ the set of vectors of size n whose coefficients are polynomials of $\mathbb{T} \mod (X^N + 1)$. N is usually a power of 2. Let $\mathbb{B} = \{0,1\}$. $\langle \, , \, \rangle$ denotes the inner product.

2.2   TFHE Scheme

The TFHE scheme is a fully homomorphic encryption scheme introduced in 2016 in [9] and implemented as the TFHE library [2]. TFHE defines three structures to encrypt plaintexts, which we summarize below as fresh encryptions of 0:

- **TLWE sample**: A pair $(a,b) \in \mathbb{T}^{n+1}$, where a is uniformly sampled from $\mathbb{T}^n$ and $b = \langle a,s \rangle + e$. The secret key s is uniformly sampled from $\mathbb{B}^n$, and the error $e \in \mathbb{T}$ is sampled from a Gaussian distribution with mean 0 and standard deviation σ.
- **TRLWE sample**: A pair $(a,b) \in \mathbb{T}_N[X]^{k+1}$, where a is uniformly sampled from $\mathbb{T}_N[X]^k$ and $b = \langle a,s \rangle + e$. The secret key s is uniformly sampled from $\mathbb{B}_N[X]^k$, the error $e \in \mathbb{T}$ is a polynomial with random coefficients sampled from a Gaussian distribution with mean 0 and standard deviation σ. One usually chooses $k = 1$; therefore, a and b are vectors of size 1 whose coefficient is a polynomial.
- **TRGSW sample**: a vector of l TRLWE fresh samples.

Let $\mathcal{M}$ denote the discrete message space ($\mathcal{M} \in \mathbb{T}_N[X]$ or $\mathcal{M} \in \mathbb{T}$). To encrypt a message $m \in \mathcal{M}$, we add what is called a *noiseless trivial* ciphertext $(0,m)$ to a fresh encryption of 0. We denote by $c = (a,b) + (0,m) = (a,b+m) \in \text{T(R)LWE}_s(m)$ the T(R)LWE encryption of m with key s. A message $m \in \mathbb{Z}[X]$ can also be encrypted in TRGSW samples by adding $m \cdot H$ to a TRGSW sample of 0, where H is a gadget decomposition matrix. As we will not explicitly need such an operation in this paper, more details about TRGSW can be found in [9].

---

[2] https://tfhe.github.io/tfhe/

To decrypt a ciphertext c, we first calculate its phase $\phi(c) = b - \langle a, s \rangle = m + e$. Then, we need to remove the error, which is achieved by rounding the phase to the nearest valid value in $\mathcal{M}$. This procedure fails if the error is greater than half the distance between two elements of $\mathcal{M}$.

2.3   TFHE Bootstrapping

Bootstrapping is the operation that reduces the noise of a ciphertext thus allowing further homomorphic calculations. It relies on three basic operations, which we briefly review in this section.

- **BlindRotate**: rotates a polynomial encrypted as a TRLWE ciphertext by an encrypted index. It takes several inputs: a ciphertext $c \in \mathrm{TRLWE}_k(m)$, a vector $(a_1, \cdots, a_p, b)$ where $\forall i, a_i \in \mathbb{Z}_{2N}$, and a TRGSW ciphertext encrypting the secret key $s = (s_1, \cdots, s_p)$.
  It returns a ciphertext $c' \in \mathrm{TRLWE}_k(m \cdot X^{\langle a, s \rangle - b})$. This paper will refer to this algorithm as `BlindRotate`.
- **TLWE Sample Extract**: extracts a coefficient of a TRLWE ciphertext and converts it into a TLWE ciphertext. It takes as inputs both a ciphertext $c \in \mathrm{TRLWE}_k(m)$ and an index $p \in \{0, \cdots, N-1\}$. The result is a TLWE ciphertext $c' \in \mathrm{TLWE}_k(m_i)$ where m_i is the i^{th} coefficient of the polynomial m. This paper will refer to this algorithm as `SampleExtract`.
- **Public Functional Keyswitching**: allows the switching of keys and parameters from p ciphertexts $c_i \in \mathrm{TLWE}_k(m_i)$ to one $c' \in \mathrm{T(R)LWE}_s(f(m_1, \cdots, m_p))$ where f is a public linear morphism between $\mathbb{T}^p$ and $\mathbb{T}_N[X]$. That is to say, this operation not only allows the packing of TLWE ciphertexts in a TRLWE ciphertext, but it can also evaluate a linear function f over the input TLWEs. This paper will refer to this algorithm as `KeySwitch`.

It is important to note that, during a `BlindRotate` operation, an excessive noise level in the input ciphertext can lead to errors in the bootstrapping output resulting in incorrect ciphertexts (i.e., ciphertext which does not decrypt to correct calculation results). This has implications for parameters and data representations choices (number of digits and basis).

Algorithm 1 shows the *TFHE Gate Bootstrapping* [9], which aims to evaluate a binary gate operation homomorphically and reduce the output ciphertext noise at the same time. To that end, 0 and 1 are respectively encoded as 0 and $\frac{1}{2}$ over $\mathbb{T}$. The first step of this algorithm consists in selecting a value $\hat{m} \in \mathbb{T}$ which will be used afterward to compute the coefficients of the polynomial which will rotate during the `BlindRotate`. We call this polynomial *testv* as seen in Step 3. Note that for any $p \in [\![0, 2N]\!]$ (where $[\![0, 2N]\!]$ corresponds to the set of integers $\{0, \cdots, 2N\}$), the constant term of $testv \cdot X^p$ is $\hat{m}$ if $p \in]\!]\frac{N}{2}, \frac{3N}{2}]\!]$ and $-\hat{m}$ otherwise. Step 4 returns an accumulator $ACC \in \mathrm{TRLWE}_{s'}(testv \cdot X^{\langle \overline{a}, s \rangle - b})$. Indeed, the constant term of ACC is $-\hat{m}$ if c is an encryption of 0 and $\hat{m}$ if c is an encryption of $\frac{1}{2}$. Then step 5 creates a new ciphertext $\overline{c}$ by extracting the constant term in position 0 from ACC and adding $(0, \hat{m})$. Thus, $\overline{c}$ corresponds to an encryption of 0 if c is an encryption of 0 and

m otherwise. On the other hand, if c is an encryption of $\frac{1}{2}$ and if we choose $m = \frac{1}{2}$, the algorithm returns a *fresh ciphertext* of $\frac{1}{2}$, that is to say the encoding of 1.

In Fig. 1, we present an example of TFHE gate bootstrapping algorithm with $\mathbb{Z}_4 = \{0,1,2,3\}$ as input space. The outer circle in Fig. 1 corresponds to the plaintext encoding in $\mathbb{T}$ as $\{0, \frac{1}{4}, \frac{2}{4}, \frac{3}{4}\}$. Meanwhile, the inner circle sets the coefficients of the test polynomial $testv$ to 1, i.e., $\hat{m} = \frac{1}{4}$. Then, we rotate the test polynomial during the bootstrapping by the phase $\phi(c_0)$ of the input ciphertext c_0. In our example, we obtain as bootstrapping output either an encryption of the encoding of 1 for positive inputs $\{0, \frac{1}{4}\}$, or an encryption of the encoding of -1 for negative inputs $\{\frac{2}{4}, \frac{3}{4}\}$.

Algorithm 1 TFHE gate bootstrapping [9]

Require: a constant $m \in \mathbb{T}$, a TLWE sample $c = (a,b) \in \text{TLWE}_s(x \cdot \frac{1}{2})$ with $x \in \mathbb{B}$, a bootstrapping key $BK_{s \to s'} = (BK_i \in \text{TRGSW}_{S'}(s_i))_{i \in [\![1,n]\!]}$ where S' is the TRLWE interpretation of a secret key s'.
Ensure: a TLWE sample $\bar{c} \in \text{TLWE}_s(x.m)$
 1: Let $\hat{m} = \frac{1}{2}m \in \mathbb{T}$ (pick one of the two possible values)
 2: Let $\bar{b} = \lfloor 2Nb \rceil$ and $\overline{a_i} = \lfloor 2Na_i \rceil \in \mathbb{Z}, \forall i \in [\![1,n]\!]$
 3: Let $testv := (1 + X + \cdots + X^{N-1}) \cdot X^{\frac{N}{2}} \cdot \hat{m} \in \mathbb{T}_N[X]$
 4: $ACC \leftarrow \texttt{BlindRotate}((0,testv),(\overline{a}_1,...,\overline{a}_n,\bar{b}),(BK_1,...,BK_n))$
 5: $\bar{c} = (0,\hat{m}) + \texttt{SampleExtract}(ACC)$
 6: return $\texttt{KeySwitch}_{s' \to s}(\bar{c})$

2.4   TFHE Functional Bootstrapping

We can use Look-Up Tables to compute functions during the bootstrapping operation. To do so, we replace the coefficients of the test polynomial $testv$ with the corresponding values of the LUT. Let us assume that we want to evaluate the function $f_\mathbb{T}$ via a LUT. Then, if we retrieve the i^{th} coefficient of $testv$, we actually get $f_\mathbb{T}(m_i)$ where m_i is the encrypted input to the bootstrapping. We refer to this idea by *programmable* or *functional* bootstrapping [10, 18, 23, 11, 12].

In Fig. 1, we give an example of functional bootstrapping with $\mathbb{Z}_4 = \{0,1,2,3\}$ as input space. We encode the images of $\{0, \frac{1}{4}\}$ by $f_\mathbb{T}$ as coefficients of the test polynomial (in the inner circle). Meanwhile, we deduce the images of $\{\frac{2}{4}, \frac{3}{4}\}$ by negacyclicity. Indeed, in $\mathbb{T}$, we can encode negacyclic functions, i.e., antiperiodic functions with period $\frac{1}{2}$ (verifying $f_\mathbb{T}(x) = -f_\mathbb{T}(x + \frac{1}{2})$), where $[0,0.5[$ corresponds to positive values and $[0.5,1[$ to negative ones. In our example, if we encrypt one of the following values $\{0, \frac{1}{4}, \frac{2}{4}, \frac{3}{4}\}$ and we give it as input to the functional bootstrapping algorithm, we get $\{f_\mathbb{T}(0), f_\mathbb{T}(1), -f_\mathbb{T}(0), -f_\mathbb{T}(1)\}$, respectively.

Almost all of the functional bootstrapping methods from state of the art ([10, 18, 23, 11, 12]) take as input a single ciphertext. In 2021, Guimarães *et al.*, [17] discussed two methods for performing functional bootstrapping with larger plaintexts. They combine several bootstrappings with different encrypted inputs by using a tree or

Fig. 1: TFHE Bootstrapping examples: the outer circles describe the inputs to the bootstrapping (i.e., ciphertexts over $\mathbb{T}$). Meanwhile, the inner circles represent the coefficients of the test polynomial *testv*. One of these coefficients is extracted as the output of the bootstrapping after the `BlindRotate`.

a chain structure. The ciphertexts are encryptions of digits that come from the decomposition of plaintexts in a certain basis B.

2.5   Tree-based Method

Let $B,B',d \in \mathbb{N}^*$ and m be an integer message. B and B' are the basis on which to decompose the message. We then have $m = \sum_{i=0}^{d-1} m_i B^i$, with $m_i \in [\![0, B-1]\!]$. From this decomposition, we obtain d TLWE encryptions $(c_0, c_1, \cdots, c_{d-1})$ of $(m_0, m_1, \cdots, m_{d-1})$ on half of the torus $\mathbb{T}$. We denote $f : [\![0, B-1]\!]^d \to [\![0, B'-1]\!]$ the target function and define g as the following bijection:

$$g : \begin{array}{ccc} [\![0, B-1]\!]^d & \to & [\![0, B^d-1]\!] \\ (a_0, a_1, \cdots, a_{d-1}) & \mapsto & \sum_{i=0}^{d-1} a_i \cdot B^i \end{array}$$

We then encode a LUT for f under the form of B^{d-1} TRLWE ciphertexts. Each of these ciphertexts encodes a polynomial P_i such that:

$$P_i(X) = \sum_{j=0}^{B-1} \sum_{k=0}^{\frac{N}{B}-1} f \circ g^{-1}(j \cdot B^{d-1} + i) \cdot X^{j \cdot \frac{N}{B} + k}$$

Then, we apply the `BlindRotateAndExtract` (the `BlindRotate` directly followed by the `SampleExtract` in position 0) to each test polynomial $testv = \text{TRLWE}(P_i)$ with c_0 as a selector. We obtain B^{d-1} TLWE ciphertexts, each corresponding to the encryption of $f \circ g^{-1}(m_{d-1} \cdot B^{d-1} + i)$, for $i \in [\![0, B^{d-1}-1]\!]$.

Finally, we use the `KeySwitch` operation from TLWE to TRLWE to gather them into B^{d-2} encrypted TRLWE, corresponding to the LUT of h, with:

$$h: \quad [\![0,B-1]\!]^{d-1} \quad \rightarrow \quad [\![0,B'-1]\!]$$
$$(a_0, a_1, \cdots, a_{d-2}) \mapsto f(a_0, a_1, \cdots, a_{d-2}, m_{d-1})$$

We then repeat this operation, using the ciphertext c_i at step i, until we obtain a single TLWE ciphertext of $f(m_0, m_1, \cdots, m_{d-1})$. Note that the tree-based method must be run independently as many times as the number of digits in the output.

2.6   Multi-value Bootstrapping

Multi-Value Bootstrapping (MVB) [8] refers to the method for evaluating k different LUTs on a single input with a single bootstrapping. MVB factors the test polynomial P_{f_i} associated with the function f_i into a product of two polynomials v_0 and v_i, where v_0 is a common factor to all P_{f_i}. In practice, we have:

$$(1 + X + \cdots + X^{N-1}) \cdot (1 - X) \equiv 2 \bmod (X^N + 1)$$

Now by writing P_{f_i} in the form $P_{f_i} = \sum_{j=0}^{N-1} \alpha_{i,j} X^j$ with $\alpha_{i,j} \in \mathbb{Z}$, we get from the previous equation:

$$P_{f_i} = \frac{1}{2} \cdot (1 + X + \cdots + X^{N-1}) \cdot (1 - X) \cdot P_{f_i} \bmod (X^N + 1)$$
$$= v_0 \cdot v_i \bmod (X^N + 1)$$

where:

$$v_0 = \frac{1}{2} \cdot (1 + X + \cdots + X^{N-1})$$
$$v_i = \alpha_{i,0} + \alpha_{i,N-1} + (\alpha_{i,1} - \alpha_{i,0}) \cdot X + \cdots + (\alpha_{i,N-1} - \alpha_{i,N-2}) \cdot X^{N-1}.$$

This factorization makes it possible to compute many LUTs using a unique bootstrapping. Indeed, it is enough to initialize the test polynomial $testv$ with the value of v_0 during bootstrapping. Then, after the `BlindRotate` operation, one has to multiply the obtained ACC by each v_i corresponding to the LUT of f_i to get ACC_i.

This optimization reduces the number of bootstrapping required for an operation and, thus, the overall computation time.

3 A Short Reminder on the AES

Advanced Encryption Standard (AES) is the name given to the Rijndael algorithm, the winner of the NIST standardization competition in 2000 [14]. It is a symmetric block cipher, defined to work with different key sizes. Several rounds are applied to the original message to obtain an encrypted message. Each round consists of the same operations performed in the same order. We chose to work on the 128-bits AES, which uses 10 rounds. The ciphertext is composed of 16 bytes such as $c = c_0 c_1 \cdots c_{15} \in (\mathbb{F}_{2^8})^{16}$ and is encoded in what we call a *state matrix* in the following way:

$$\begin{pmatrix} c_0 & c_4 & c_8 & c_{12} \\ c_1 & c_5 & c_9 & c_{13} \\ c_2 & c_6 & c_{10} & c_{14} \\ c_3 & c_7 & c_{11} & c_{15} \end{pmatrix}$$

The round operations affect this matrix as follows:

- `SubBytes`: the `SubBytes` operation is the only non-linear transformation of the cipher. It is a permutation consisting of an S-box applied to the bytes of the state matrix. As it acts on the individual bytes of the state, it can be parallelized for efficient execution.
- `AddRoundKey`: before the encryption, the secret key is "expanded" into several round keys. The encryption process relies only on these round keys and not the initial secret key. In this transformation, the state is modified by combining it with a round key with the bitwise XOR operation. Of course, to do so, the size of each round key is equal to the size of the ciphertext (in our case, 128 bits encoded in a round key matrix to match the state matrix). That is to say; we have at round $i \in \{0, \cdots, 9\}$:

$$\begin{pmatrix} c_0 & c_4 & c_8 & c_{12} \\ c_1 & c_5 & c_9 & c_{13} \\ c_2 & c_6 & c_{10} & c_{14} \\ c_3 & c_7 & c_{11} & c_{15} \end{pmatrix} \oplus \begin{pmatrix} k_{i,0} & k_{i,4} & k_{i,8} & k_{i,12} \\ k_{i,1} & k_{i,5} & k_{i,9} & k_{i,13} \\ k_{i,2} & k_{i,6} & k_{i,10} & k_{i,14} \\ k_{i,3} & k_{i,7} & k_{i,11} & k_{i,15} \end{pmatrix} = \begin{pmatrix} c_0 \oplus k_{i,0} & c_4 \oplus k_{i,4} & c_8 \oplus k_{i,8} & c_{12} \oplus k_{i,12} \\ c_1 \oplus k_{i,1} & c_5 \oplus k_{i,5} & c_9 \oplus k_{i,9} & c_{13} \oplus k_{i,13} \\ c_2 \oplus k_{i,2} & c_6 \oplus k_{i,6} & c_{10} \oplus k_{i,10} & c_{14} \oplus k_{i,14} \\ c_3 \oplus k_{i,3} & c_7 \oplus k_{i,7} & c_{11} \oplus k_{i,11} & c_{15} \oplus k_{i,15} \end{pmatrix}$$

- `ShiftRows`: the `ShiftRows` step is a byte transposition that cyclically shifts the rows of the state over different offsets. For AES-128, row 0 is shifted over 0 bytes, row 1 over 1 byte, row 2 over 2 bytes and row 3 over 3 bytes. As this operation only alters the position of the bytes in the state matrix, it does not require a homomorphic equivalent (instead of shifting regular bytes, we shift homomorphic bytes).
- `MixColumns`: the `MixColumns` step is operating on the state column by column via matrix multiplication. But in practice, the authors do not implement the naive matrix product but work with each byte of the state matrix individually. They use a mix of scalar multiplication and XOR, as visible in the original code. Therefore, this operation can be parallelized.

A typical execution of the 128-bits AES begins with the first `AddRoundKey` directly followed by the first iteration of the rounds. Each round proceeds as follows:

1. `SubBytes`

2. `ShiftRows`
3. `MixColumns`
4. `AddRoundKey`

Except for the last round, which does not require the `MixColumns` step.

About the Key expansion– `Key expansion` is an operation that may be performed once and for all, from the secret key. Indeed, from the 128-bit key are derived eleven 128-bit round keys, which are used in the `AddRoundKey` operation. As a result, to evaluate the AES encryption or decryption algorithm, a server only needs to know the round keys. This operation, consisting of XOR and multiplication, is expensive in the homomorphic domain. It is therefore more efficient for a client to generate its own key, derive the round keys, and then homomorphically encrypt them. Sending these eleven homomorphically encrypted round keys is faster than creating the initial key, encrypting it in homomorphic, and sending it to a server to perform the `Key Schedule` in the homomorphic domain. For these reasons, we remove the key schedule from the encrypted-domain computations.

4   AES goes homomorphic

To use the full potential of programmable bootstrapping, we aim to transform the AES algorithm into a succession of LUT evaluations. The permutation given by the Sbox is already in LUT form, so all that remains is to modify the multiplication in $GF(256)$ and XOR operations. These are binary operations, which are not immediate to put into LUT form. Indeed, evaluating a LUT is like dereferencing an array from an encrypted index. The LUT evaluation operator, therefore, takes only one input: this encrypted index. If possible, these operations should thus be transformed into unary table indirection operations. Since AES works on a byte-by-byte basis, LUTs must take 8-bits encrypted indexes and return an evaluation of this index on 8 bits. In this section, we explain the various steps involved in transforming the original code into the sequence of 8-bits-to-8-bits LUTs required for efficient homomorphic execution.

4.1   Optimizing the Original AES Code

The first step towards the homomorphization of a symmetric cryptosystem is to look at the original code for small changes that would ease the transition. AES makes no exception. However, we quickly observed that the proposed C++ implementation (given by the creators of the AES in [14]) is only "pseudo" 8-bits, as a carry that requires an additional ninth bit is necessary for specific bytes operations.

Indeed, the problem occurs during the multiplication of two integers (see Fig. 2). To compute this operation efficiently, the authors use a generator $\alpha = X + 1$ of the message space $GF(256) \cong \mathbb{F}_2[X]/\langle X^8 + X^4 + X^3 + X + 1\rangle$. From this, they construct two tables of 256 elements. The first one is `Logtable`, defined such as `Logtable`$[\alpha^i] = i$. The second is `Alogtable` where `Alogtable`$[i] = \alpha^i$. Then, instead of a naive multiplication, the result is obtained with a simple sum of logs and two table indirections.

Consequently, when the `mul`(word8 a, word8 b) function is called, two 8-bits integers representing the logs are added together. And this sum may exceed 255, which is why the authors apply a %255 to the result. We have to change the whole function structure to avoid this overhead and have an actual 8-bits implementation. But it also implies some more profound changes in the entire code.

```
word8 mul(word8 a, word8 b) {
  /* multiply two elements of GF(256)
  * required for MixColumns and InvMixColumns
  */
  if (a && b) return Alogtable[(Logtable[a] + Logtable[b])%255];
  else return 0;
}
```

Fig. 2: The sum of `Logtable[a] + Logtable[b]` may exceed 255. Therefore, this operation requires more than 8 bits in practice, which is unsuitable for a homomorphic evaluation with an 8-bits processor.

We fix this issue by observing that there are only a few calls to the `mul` function throughout the algorithm. During these calls, the parameter a only takes six different values. Indeed, a careful reading of the code shows that $a \in \{2, 3, 9, b, d, e\}$. This means that we can separate the `mul`(a,b) function into six `mul_a`(b) functions, with $a \in \{2,3,9,b,d,e\}$ by removing the variable a of the parameters and just encoding the correct values in the corresponding new functions. But this does not solve the issue. To suppress the problematic addition of two 8-bits integers, we compute and hard-code six tables `Tlog_a` such that `Tlog_a`$[b] = $`(Logtable[a]+Logtable[b])%255`. This pre-calculation of the tables enables us to eliminate the problem of the "pseudo" 8-bits implementation. At the same time, this allows us to put an end to the modulo operation, which does not really have a known operator for this use in FHE. To take optimization a step further and reduce the number of tables indirections per encrypted index, we create six new `T_a` tables such that `T_a`$[b]=$`Alogtable[Tlog_a[b]]` and encode them directly into our implementation.
We now have six functions of the following form:

```
word8 mul_a(word8 b) {
    if (b) return T_a[b];
    else return 0;
}
```

It is a simple function, but regarding a homomorphic evaluation, we want our function to be as light as possible regarding the number of operations performed. This is why, to get rid of the `if` condition on b, we modified the `T_a` tables so that they consider the case where $b=0$. That is to say, for every $a \in \{2,3,9,b,d,e\}$, `T_a`$[0]=0$. So our final implementations of the multiplications are straightforward and only require one indirection, as seen in Fig. 3.

```
word8 mul_a(word8 b) {
  return T_a[b];
}
```

Fig. 3: The new multiplication functions are very simple and easy to execute in the homomorphic domain, thanks to pre-calculation and tables hard-encoding.

Therefore, this work on the multiplication functions allows us to operate with only one indirection instead of three indirections, an addition, a modulo operation, and an if condition in the original code version. These changes allow us to convert the initial binary multiplication into several unary multiplications, which is ideal for a LUT transformation.

When looking for other possible optimizations, we realized we achieved an optimal or nearly optimal form regarding LUT factorization and, thus, LUT-based homomorphic evaluation (at least when starting from the standard AES implementation). Indeed, the structure of the `MixColumns` instructions prevent further factorisation. For example, if one consider the following instruction (in `MixColumn`)

```
b[0][j] = mul2(a[0][j]) ^ mul3(a[1][j]) ^ a[2][j] ^ a[3][j];
```

then, although the `mul2` and `mul3` LUT could directly embed the Sbox LUT (assuming the indices are made consistent with the effect of `ShiftRows`) the two last terms cannot (since they do not require any $GF(256)$ multiplication, hence there is no LUT to factor the Sbox LUT with). Still, as an additional optimisation, we have also merged the `AddRoundKey` function with the `SubBytes` one, as the first one always precedes the second one, for increased parallelism efficiency.

4.2   And Finally Going Homomorphic for Real

To make the most of TFHE programmable bootstrapping, we implement all operations (XOR, Sbox, multiplication) via LUT evaluations. As we now have a ciphertext-by-cleartext $GF(256)$ multiplication function that is performed only by means of a unique indirection, this causes no issue. And by definition, the Sbox is a LUT. So the actual work here is to transform the XOR into a LUT. The XOR is a binary operation in which, in our case, inputs are bytes. That means we need a 256×256 table to encode our XOR. As it is a binary operation, we must either use the tree-base method or the chaining method, both introduced in [17], which allows bootstrapping with several entries. Table 1 compares the parameters needed depending on the method. Relying on this comparison table, we chose to work with the tree-based method.

The MVB optimization allows us to evaluate a LUT with two input digits and one output digit in only two bootstrapping. Still, as the parameters needed are significant (we have to work on the discretized torus on 512 values, using only the positive half), it will not be efficient. Indeed, we need to use the cyclotomic polynomial $X^{32768} + 1$, which results in very slow bootstrapping. This makes it interesting to break down the messages into smaller bases and to use the tree-based method for every LUT evaluation. We observe that the number N_{boot} of bootstrapping needed with the tree-based method increases with the number of input digits d and the

	decomposition basis	n	N
	4	850	2048
chaining method	8	1024	8192
	16	1100	32768
	4	700	1024
tree-based method	8	700	2048
	16	1024	2048

Table 1: Parameters needed depending on the method and the chosen basis. Given a message space of size B, the chaining method requires to use a plaintext space of size B^2 (instead of only $2 \times B$ for the tree-based method). As such, the size of the parameters dramatically increases as the basis B increases. This growth of parameters jeopardizes the other speed improvements that could come with the method. Hence we stuck to the tree-based method that we find more efficient for computing non-polynomial functions in our case.

number of output digits d' as well as the basis B chosen for the decomposition. In fact, for a complete evaluation of a LUT via the classic tree-based method, we have $N_{boot} = d' \times \sum_{i=0}^{d-1} B^i$ and with the MVB, we have $NB'_{boot} = d' \times (1 + \sum_{i=0}^{d-2} B^i)$.

   Although the tree-based method requires more operations, the smaller the decomposition basis, the smaller the parameters to be used. Small ciphertext sizes allow small parameters and, therefore, faster operations. A balance must therefore be struck between the number of bootstrapping operations to be performed and the size of the parameters. For this reason, we produce a benchmark of the execution time of LUT evaluation depending on the ciphertext decomposition.
We analyze different decompositions:

- basis 256: it is not a decomposition per see, but we have to be sure that this basis is not the most advantageous one
- basis 16: the message is decomposed into two digits in basis 16 (4 bits per digit)
- basis 8: the message is decomposed into three digits in basis 8 (3 bits per digit, but only 2 bits for the most significant digit)
- basis 4: the message is decomposed into four digits in basis 4 (2 bits per digit)

For this, we first implement an efficient homomorphic operator to evaluate any LUT using TFHE programmable bootstrapping and MVB (and the TFHE library). We then use it with our different decompositions and matching parameter sets. The result of this study is resumed in Table 2. Regarding the execution time per LUT evaluation, we choose to work with the decomposition in basis 16. This choice explains our implementation using messages in basis 16 instead of whole bytes. Using our efficient homomorphic operator, we transform the `mul_a` functions to be homomorphic. To do so, we decompose every table `T_a` into two tables in basis 16, one per digit of the decomposition. We do the same with the Sbox table. Finally, we create a table for a 4-bits XOR. This table is of size 16×16, which is easier to encode than the 256×256 needed without the decomposition. This allows us to "translate" our optimized version of the AES code into an efficient homomorphic one.

decomposition basis	single bootstrapping	#boot. for MVB	complete LUT evaluation	parameter n	parameter N
256	1.5 s	1	1.5 s	1024	32768
16	0.029 s	4	0.3 s	1024	2048
8	0.015 s	30	1.4 s	700	2048
4	0.007 s	88	2.0 s	700	1024

Table 2: Unitary timings for bootstrapping and full LUT evaluation depending on basis choices. There is no linearity between the timings due to the number of multiplications involved by the MVB.

4.3   Details about the Homomorphization of the XOR Operator

We have to transform the XOR operator. On the one hand, we have to transform it to fit the decomposition basis chosen, and on the other hand, we have to transform it so we can evaluate it via a LUT. Indeed, we chose to work with basis 16. That means that our messages are of the form $m = m_0 + m_1 \cdot 16$, and the corresponding ciphertext is a vector of the form $c = (c_0, c_1) = ([m_0], [m_1])$. So, to compute the XOR of two ciphertexts c and c' in basis 16, we have to compute:

$$c \oplus c' = (c_0, c_1) \oplus (c'_0, c'_1) = (c_0 \oplus c'_0, c_1 \oplus c'_1)$$

We thus have to evaluate a 4-bits-by-4-bits XOR on the ciphertexts. For this, we compute the double input table of 4-bits XOR, which has size 16×16. The tricky part is to choose the correct way to construct the test polynomials from this table for the tree-based method. Indeed, on the first step of the tree-based method, we have 16 test polynomials. Each one must encode the coefficient of a unary XOR. That is to say, the first polynomial P_0 encodes the values of the unary operation `xor_by_0`(i) such that for $i \in \{0, 1, \cdots, 15\}$ `xor_by_0`$(i) = i \oplus 0$, the second polynomial P_1 encodes `xor_by_1`, etc. So when applying the `BlindRotate` on these polynomials with selector $c_0 = [m_0]$, we obtain 16 new ciphertexts encoding `xor_by_m0` that we put together on a polynomial P_{final}. We then apply the `BlindRotate` on the new polynomial P_{final} with selector $c'_0 = [m'_0]$, and we obtain the final results, which is an encryption of $c_0 \oplus c'_0$. The method is resumed in Fig. 4.

By using the same method, we can easily compute $c_1 \oplus c'_1$, and thus obtain the vector $c \oplus c' = (c_0 \oplus c'_0, c_1 \oplus c'_1)$ encoding $m \oplus m'$.

4.4   Parallelization

On the one hand, the purpose of transciphering is to avoid transferring large ciphertexts. Since a server has more computing power than a client, it can efficiently run on several cores to parallelize computations and optimize execution times. On the other hand, AES, as a block cipher, is naturally parallelizable: each round operation can be performed simultaneously on each byte of the state matrix (except for the `ShiftRows` step). This is why, in continuing our work, we use the OpenMP library to parallelize our code and optimize the execution times of our homomorphic AES. We run our tests on two machines. The first one is a 12th Gen Intel(R) Core(TM)

Fig. 4: The principle of the tree-based method applied to the XOR with basis 16. The green color indicates that the contents of the box is encrypted.

i7-12700H CPU (using six cores) laptop with 64 Gib total system memory with an Ubuntu 22.04.2 LTS server. The second is an AMD EPYC 7702P 64-cores Processor server with an Ubuntu 20.04.6 LTS server. In the following sections, we will refer to them respectively as `i7-laptop` and `AMD-server`.

Using the OpenMP library, we first parallelize onto the six cores on the `i7-laptop`. As it is a standard laptop, it is interesting to see how a partially parallelized homomorphic AES can work. Indeed, it gives a good order of magnitude of the possibilities of a full parallelization. That is why we parallelize every round function, except for the `ShiftRows` one as it only involves ciphertexts reorganization within the state matrix. Still using the OpenMP library, we were able to use 16 cores for a full parallelization on the `AMD-server`. As the server is small (only 64 cores) and a bit old, it is slower than the `i7-laptop` in sequential time. But, by running a sequential execution, and a full parallelized one, we can calculate the speedup factor induced by parallelization. For instance, the execution parallelized on 16 cores is 9.5 times faster than the sequential one on the `AMD-server`. Results can be found in Table 3.

A perspective is to go further in the parallelization of the AES evaluation. Indeed, we can also parallelize the tree-based method and execute the computation of each decomposition digit at the same time. For this method, we will need 32 cores, but it can divide the execution time almost by 2.

5   Computation times

This section resumes the results of our implementations in Table 3. The speed-ups are inferior to the number of cores used due to the call to the multiple threads and their

closing at each call of functions. We improve the state of the art as the so far best-known implementation runs in 4.2 minutes with 16 threads parallelization, although no implementation details were provided (to the best of our knowledge) beyond a poster presented by Stracovsky et al., at FHE.org 2022 [22]. As Table 3 shows, we obtain a comparable *sequential* time on `i7-laptop` and a very improved one with parallel execution. Indeed, our parallelized version on `i7-laptop` is 4.6 times more efficient than Stracovsky et al.'s [22] and almost 20 times more efficient than Gentry et al.'s. [16]. Furthermore, our parallelized version on `AMD-server` is almost 7 times faster than [22], and 30 times faster than [16] as it only takes 3% of its execution time.

	seq. exec time	par. exec time	# threads	speed up
`i7-laptop`	4.5 mins = 270 secs	54.31 secs	6	4.97
`AMD-server`	5.7 mins = 342 secs	36.39 secs	16	9.4
`"i7-server"`	-	28.73 secs	16	-
Gentry et al. [16]	18 mins	-	1	-
Mella and Susella [21]	22 mins	-	1	-
Stracovsky et al. [22]	-	4.2 mins = 252 secs	16	-

Table 3: Our execution times compared to the state of the art (the `"i7-server"` time corresponds to an extrapolation of the speedup observed with 16 threads on the `AMD-server` from the `i7-laptop` sequential time as we did not have an i7-based server to perform experiments).

6   Conclusion

In this paper, we have proposed an homomorphic AES implementation relying on two "instructions", a unary 256-by-8 bits table indirection and a binary 8-bit XOR instruction, both running over encryptions of nibbles (the cute name for hex digits) and relying on functional bootstrapping for efficiency. Our work illustrates that, even when starting from the standard AES implementation, this approach significantly improves the state of the art of homomorphic AES execution timings. In terms of perspectives, beyond improved parallelism, it would be interesting to consider other non standard forms for the AES as a starting point in search for ones which may lead to smaller numbers of homomorphic operations (i.e., in fine, less bootstrappings), keeping in mind that our XOR instruction is twice as costly as the indirection one (so maintaining an appropriate balance between the two). Another interesting perspective would be to extend this embryonic "instruction set" in order to more easily apply this approach to other algorithms.

References

1. Albrecht, M., Rechberger, C., Schneider, T., Tiessen, T., Zohner, M.: Ciphers for mpc and fhe. Cryptology ePrint Archive, Paper 2016/687 (2016), https://eprint.iacr.org/2016/687, https://eprint.iacr.org/2016/687

2. Ashur, T., Mahzoun, M., Toprakhisar, D.: Chaghri - a fhe-friendly block cipher. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. p. 139–150. CCS '22, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3548606.3559364, https://doi.org/10.1145/3548606.3559364

3. Bendoukha, A.A., Boudguiga, A., Sirdey, R.: Revisiting stream-cipher-based homomorphic transciphering in the tfhe era. In: Foundations and Practice of Security: 14th International Symposium, FPS 2021, Paris, France, December 7–10, 2021, Revised Selected Papers. p. 19–33. Springer-Verlag, Berlin, Heidelberg (2021). https://doi.org/10.1007/978-3-031-08147-7_2

4. Bendoukha, A., Clet, P.E., Boudguiga, A., Sirdey, R.: Optimized stream-cipher-based transciphering by means of functional bootstrapping. DBSEC (in press) (2023), https://eprint.iacr.org/2016/687, https://eprint.iacr.org/2016/687

5. Borghoff, J., Canteaut, A., Güneysu, T., Kavun, E.B., Knezevic, M., Knudsen, L.R., Leander, G., Nikov, V., Paar, C., Rechberger, C., Rombouts, P., Thomsen, S.S., Yalçin, T.: Prince - a low-latency block cipher for pervasive computing applications - extended abstract. In: International Conference on the Theory and Application of Cryptology and Information Security (2012)

6. Cannière, C.D., Preneel, B.: Trivium. In: Robshaw, M.J.B., Billet, O. (eds.) New Stream Cipher Designs - The eSTREAM Finalists, Lecture Notes in Computer Science, vol. 4986, pp. 244–266. Springer (2008). https://doi.org/10.1007/978-3-540-68351-3_18, https://doi.org/10.1007/978-3-540-68351-3_18

7. Canteaut, A., Carpov, S., Fontaine, C., Lepoint, T., Naya-Plasencia, M., Paillier, P., Sirdey, R.: Stream ciphers: A practical solution for efficient homomorphic-ciphertext compression. Cryptology ePrint Archive, Paper 2015/113 (2015), https://eprint.iacr.org/2015/113, https://eprint.iacr.org/2015/113

8. Carpov, S., Izabachène, M., Mollimard, V.: New techniques for multi-value input homomorphic evaluation and applications. Cryptology ePrint Archive, Paper 2018/622 (2018), https://eprint.iacr.org/2018/622

9. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: Fast fully homomorphic encryption library (August 2016), https://tfhe.github.io/tfhe/

10. Chillotti, I., Joye, M., Paillier, P.: Programmable bootstrapping enables efficient homomorphic inference of deep neural networks. In: Dolev, S., Margalit, O., Pinkas, B., Schwarzmann, A. (eds.) Cyber Security Cryptography and Machine Learning. pp. 1–19. Springer International Publishing, Cham (2021)

11. Chillotti, I., Ligier, D., Orfila, J.B., Tap, S.: Improved programmable bootstrapping with larger precision and efficient arithmetic circuits for tfhe. Cryptology ePrint Archive, Report 2021/729 (2021), https://ia.cr/2021/729

12. Clet, P.E., Zuber, M., Boudguiga, A., Sirdey, R., Gouy-Pailler, C.: Putting up the swiss army knife of homomorphic calculations by means of tfhe functional bootstrapping. Cryptology ePrint Archive, Paper 2022/149 (2022), https://eprint.iacr.org/2022/149

13. Cosseron, O., Hoffmann, C., Méaux, P., Standaert, F.X.: Towards globally optimized hybrid homomorphic encryption - featuring the elisabeth stream cipher. Cryptology ePrint Archive, Paper 2022/180 (2022), https://eprint.iacr.org/2022/180, https://eprint.iacr.org/2022/180

14. Daemen, J., Rijmen, V.: The Design of Rijndael: AES - The Advanced Encryption Standard (Information Security and Cryptography). Springer, 1 edn. (2002)

15. Dobraunig, C., Grassi, L., Helminger, L., Rechberger, C., Schofnegger, M., Walch, R.: Pasta: A case for hybrid homomorphic encryption. IACR Cryptol. ePrint Arch. **2021**, 731 (2021)

16. Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the aes circuit. In: Safavi-Naini, R., Canetti, R. (eds.) Advances in Cryptology – CRYPTO 2012. pp. 850–867. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
17. Guimarães, A., Borin, E., Aranha, D.F.: Revisiting the functional bootstrap in tfhe. IACR Transactions on Cryptographic Hardware and Embedded Systems **2021**(2), 229–253 (Feb 2021). https://doi.org/10.46586/tches.v2021.i2.229-253
18. Kluczniak, K., Schild, L.: FDFB: Full domain functional bootstrapping towards practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2021/1135 (2021), https://ia.cr/2021/1135
19. Lepoint, T., Naehrig, M.: A comparison of the homomorphic encryption schemes fv and yashe. In: Pointcheval, D., Vergnaud, D. (eds.) Progress in Cryptology – AFRICACRYPT 2014. pp. 318–335. Springer International Publishing, Cham (2014)
20. Liu, F., Anand, R., Wang, L., Meier, W., Isobe, T.: Coefficient grouping: Breaking chaghri and more. In: Hazay, C., Stam, M. (eds.) Advances in Cryptology – EUROCRYPT 2023. pp. 287–317. Springer Nature Switzerland, Cham (2023)
21. Mella, S., Susella, R.: On the homomorphic computation of symmetric cryptographic primitives. In: Stam, M. (ed.) Cryptography and Coding. pp. 28–44. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
22. Stracovsky, R., Mahdavi, R.A., Kerschbaum, F.: Faster evaluation of aes using tfhe. Poster Session, FHE.Org (2022)
23. Yang, Z., Xie, X., Shen, H., Chen, S., Zhou, J.: Tota: Fully homomorphic encryption with smaller parameters and stronger security. Cryptology ePrint Archive, Report 2021/1347 (2021), https://ia.cr/2021/1347