

Optimized Homomorphic Evaluation of Boolean Functions

Nicolas Bon^{1,2}, David Pointcheval², and Matthieu Rivain¹

¹ CryptoExperts

`firstname.lastname@cryptoexperts.com`

² Ecole Normale Supérieure - PSL

`firstname.lastname@ens.fr`

Abstract. We propose a new framework to homomorphically evaluate Boolean functions using the Torus Fully Homomorphic Encryption (TFHE) scheme. Compared to previous approaches focusing on Boolean gates, our technique can evaluate more complex Boolean functions with several inputs using a single bootstrapping. This allows us to greatly reduce the number of bootstrapping operations necessary to evaluate a Boolean circuit compared to previous works, thus achieving significant improvements in terms of performances. We define theoretically our approach which consists in adding an intermediate homomorphic layer between the plain Boolean space and the ciphertext space. This layer relies on so-called p -encodings embedding bits into \mathbb{Z}_p . We analyze the properties of these encodings to enable the evaluation of a given Boolean function and provides a deterministic algorithm (as well as an efficient heuristic) to find valid sets of encodings for a given function. We also propose a method to decompose any Boolean circuit into Boolean functions which are efficiently evaluable using our approach. We apply our framework to homomorphically evaluate various cryptographic primitives, and in particular the AES cipher. Our implementation results show significant improvements compared to the state of the art.

1 Introduction

Homomorphic encryption (HE) is a cryptographic technique allowing the computation of operations on encrypted messages (which directly reflect on the original messages once decrypted), using only knowledge of public data. For example, an additive homomorphic encryption scheme is able to encrypt two messages m_1 and m_2 in ciphertexts c_1 and c_2 and to compute a third ciphertext c_3 from c_1 and c_2 that encrypts the sum $m_1 + m_2$, without knowledge of the secret key.

The security of these schemes relies on a small *noise* introduced in the data when encrypting. The problem arising is that this noise is growing while homomorphic computations are carried out, which bury the original data into the noise and makes it unrecoverable at decryption. In 2009, Gentry [18] introduced the operation of *bootstrapping* to solve this problem. This operation resets the noise at a nominal level *without decryption* allowing a potentially infinite amount of operations, making the construction of a scheme achieving *Fully Homomorphic Encryption* (FHE) possible. This operation being extremely heavy and slow, it is considered as the main bottleneck for the development of schemes efficient enough to be used in practice.

Currently, the most popular schemes in the FHE ecosystem are lattice-based and rely on the hardness of the Learning With Errors assumption [24] and/or its ring variant RLWE [22]. BFV [7], BGV [8] and CKKS [12] are *leveled* schemes, which means that they keep track of the “level” of noise in the data during the homomorphic evaluation. As soon as this level reaches a critical bound, no more computations can be performed. Some recent works (see e.g [11], [10], [21]) propose a bootstrapping operation for these schemes but it is not considered as efficient enough yet. On the other hand, TFHE [13] is built on top of a powerful bootstrapping technique known to currently be the most efficient but limiting the precision of encrypted data.

Each FHE scheme offers a set of basic homomorphic operations that can be used to build more complex algorithms. In general, these operations are homomorphic additions and multiplications, however some complex operations cannot be constructed only with these operations. TFHE offers homomorphic additions and multiplications by a plaintext as well, but its force lies in its operation of *programmable bootstrapping* allowing

the evaluation of encrypted look-up tables (LUT) while resetting the noise level. However, for performance issues, these look-up tables can only handle a small amount of bits as input (around 8 bits maximum) so the scheme is best suited for applications requiring a small precision.

In particular, TFHE is the best option to evaluate Boolean circuits with encrypted inputs, but the performances of the existing frameworks are still limited. In [13], the authors propose a strategy to evaluate Boolean functions called the *gate bootstrapping*, in which they perform one bootstrapping for each bivariate Boolean gate of the underlying circuit. As a consequence, the conversion of the original Boolean circuit in a homomorphic circuit handling encrypted bits is straightforward, moreover the noise growth is contained thanks to the systematic use of bootstrapping. However, this approach is very expensive due to the high numbers of bootstrappings and makes it highly suboptimal for large circuits.

The authors of [14] propose a different approach: by leveraging a newer version of the TFHE scheme supporting a new operation named *TLWE ciphertexts multiplication*, Boolean circuits are evaluated with homomorphic sums for XOR gates and this new multiplication operation for AND gates. If this approach is clearly a progress from the vanilla framework, we note that a few bootstrappings are still required to control the noise growth and that this new operation of TLWE multiplications remains costly both in terms of performances and in terms of noise. Thus, we choose to stick to the first version of the TFHE scheme (while slightly modifying it) to keep the framework lighter and we tackle the performance issues of [13] with a different approach than the one of [14].

Our work introduces a new framework to homomorphically evaluate Boolean functions on encrypted data efficiently, i.e. by reducing the amount of necessary bootstrappings. Our approach introduces an *intermediate homomorphic layer* which encodes bits on a small ring \mathbb{Z}_p before encrypted them. This allows us to evaluate Boolean functions with one cheap homomorphic sum followed by one bootstrapping. After formalizing the underlying concept of *p-encoding* and explaining our evaluation strategy, we investigate the issue of finding valid sets of encodings for a Boolean function. We characterize this problem and provide an exact constructive algorithm to solve it. We further provide a sieving heuristic finding solutions more efficiently but at the cost of loosing optimally. Since our method is only efficient for Boolean functions with limited number of inputs, we also propose a heuristic to decompose any Boolean circuit into Boolean functions which are efficiently evaluable using our approach. Finally, we apply our technique to various cryptographic primitives, namely the SIMON block cipher, the Trivium stream cipher, the Keccak permutation, the Ascon s-box and the AES s-box. Compared to previous works implementing the same primitives (for SIMON, Trivium and AES) our implementations achieve significant speedups.

After some technical preliminaries on TFHE (Section 2), we introduce a new concept of *intermediate homomorphic layer* and explain how bits are encoded in Section 3. and the algorithms to construct it in Sections 4, 5. Finally, we describe our modifications of the TFHE scheme in Section 6 and our experimental results in Section 7.

2 Preliminaries on TFHE

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. In practice, torus elements are not represented with an infinite number of digits. Let us denote this precision in base 2 as Ω . We can define the discretized torus $\mathbb{T}_q = \{\frac{a}{q} \mid a \in \mathbb{Z}_q\}$ (the elements of the torus up to Ω bits of precision, q being 2^Ω) and identify it with the ring \mathbb{Z}_q . As a consequence, any element $\frac{a}{q}$ of \mathbb{T}_q will be represented in machine by a without any loss of property of the group \mathbb{T}_q . The operations of sum $+$ and external product \cdot have to be understood modulo q .

Moreover, for a natural integer N and a given q , we will denote by $\mathbb{T}_{N,q}[X]$ the ring of polynomial $\mathbb{T}_q[X]/(X^N + 1)$. The elements of this ring are polynomials of maximum degree $N - 1$ and with coefficients in \mathbb{T}_q . Like for the scalar version, this ring will be identified with the ring $\mathbb{Z}_q/(X^N + 1)$. N is usually taken as a power of two.

Finally, we will denote by \mathbb{B} the set of binary digits $\{0, 1\}$. $\&$ and \oplus denote the AND and XOR binary operations. For x and $z \in \mathbb{Z}$, $[x]_q$ denotes the reduction of x modulo q . For S a set, $x \stackrel{\$}{\leftarrow} S$ denotes a uniformly random sampling from the set. For χ a distribution, $x \stackrel{\$}{\leftarrow} \chi$ denotes a random sampling according to the distribution.

2.2 Complexity Assumptions

The TFHE scheme, as other schemes using lattices, relies on the hardness of the LWE assumption. More precisely, it relies on the torus-based version of the problem. In the following, we consider the classic definition but over a discretized torus and with a binary secret:

Definition 1. (*LWE problem over the discretized torus*). Let $q, n \in \mathbb{N}$ and let $\mathbf{s} = (s_1, \dots, s_n) \stackrel{\$}{\leftarrow} \mathbb{B}^n$. Let χ be an error distribution over \mathbb{Z}_q . The decisional Learning With Errors over discretized torus problem is to distinguish samples chosen with the following distributions:

$$\mathcal{D}_0 = \{(\mathbf{a}, r) \mid \mathbf{a} \stackrel{\$}{\leftarrow} \mathbb{T}_q^n, r \stackrel{\$}{\leftarrow} \mathbb{T}_q\}$$

and:

$$\mathcal{D}_1 = \{(\mathbf{a}, b) \mid \mathbf{a} = (a_1, \dots, a_n) \stackrel{\$}{\leftarrow} \mathbb{T}_q^n, e \stackrel{\$}{\leftarrow} \chi, b = \sum_{j=1}^n a_j \cdot s_j + e\}$$

The search version of the problem is to recover \mathbf{s} from the samples of \mathcal{D}_1 .

Both the search and decisional problems are reducible to each other [24] and their average case is as hard as worst-case lattice problems.

[20] argues that identifying the discretized torus \mathbb{T}_q as \mathbb{Z}_q makes the LWE assumption over the discretized torus as hard as the standard LWE assumption.

TFHE relies as well on the generalized version of LWE over rings introduced in [8] named GLWE.

Definition 2. (*GLWE problem over the discretized torus*). Let $N, q, k \in \mathbb{N}$ with N a power of two and let $\mathbf{s} = (s_1, \dots, s_k) \stackrel{\$}{\leftarrow} \mathbb{B}_N[X]^k$. Let χ be an error distribution over $\mathbb{Z}_{N,q}[X]$. The General decisional Learning With Errors over discretized torus problem is to distinguish samples chosen with the following distributions:

$$\mathcal{D}_0 = \{(\mathbf{a}, r) \mid \mathbf{a} \stackrel{\$}{\leftarrow} \mathbb{T}_{N,q}[X]^k, r \stackrel{\$}{\leftarrow} \mathbb{T}_{N,q}[X]\}$$

and:

$$\mathcal{D}_1 = \{(\mathbf{a}, b) \mid \mathbf{a} = (a_1, \dots, a_n) \stackrel{\$}{\leftarrow} \mathbb{T}_{N,q}[X]^k, e \stackrel{\$}{\leftarrow} \chi, b = \sum_{j=1}^k a_j \cdot s_j + e\}$$

The search version is analogous to the LWE one.

Note that RLWE is simply an instantiation of GLWE with $k = 1$.

The complexity analysis is analogous to the LWE version. In practice, the error distribution χ is a centered Gaussian distribution parametrized by its standard deviation σ .

2.3 Plaintext Space

Before expliciting more in depth the TFHE scheme, it is useful to define the *plaintext space* and how it is embedded in the discretized torus.

The plaintext space is the ring \mathbb{Z}_p , with $p \in \mathbb{N}$. For now, let us assume that $p \mid q$ and identify \mathbb{Z}_p with \mathbb{T}_p . As $p \mid q$, all elements of \mathbb{T}_p are elements of \mathbb{T}_q as well. Thus, we can define a mapping $\rho : \mathbb{Z}_p \rightarrow \mathbb{Z}_q$

as $\rho : m \mapsto \frac{mq}{p}$. Of course, only p elements of \mathbb{Z}_q are reached by such a mapping and they have the form $\left\{ \frac{kq}{p} \mid k \in \mathbb{Z}_p \right\}$. As they are evenly distributed across \mathbb{Z}_q , they define what we call *sectors of \mathbb{Z}_q* of the form:

$$\left\{ \left(\frac{(2k-1)q}{2p}, \frac{(2k+1)q}{2p} \right) \mid k \in \mathbb{Z}_p \right\}.$$

During encryption of m , some small noise e is drawn from a Gaussian distribution over \mathbb{Z}_q and is added to m . As e is small, the noisy message $m+e$ stays in the same sector as m but while homomorphic operations are carried out, the noise grows and may overflow out of the sector. When decrypting, one recovers the sum of the expected result and some noise $m'+e'$. As long as $e' < \frac{q}{2p}$, the message m' can be recovered by rounding to the closest center of sector.

In our work, we pick some values of p that does not divide q . Consequently, the centers and the bounds of sectors are computed by rounding the fractions to the closest integers. In practice, p is much smaller than q (p is restricted to a few bits, while q typically equals 2^{32} or 2^{64}), so this discrepancy makes this approximation sound. In the following, we will ignore this rounding.

2.4 Ciphertexts Types and Basic Operations

TFHE manipulates several different types of ciphertexts. In the following, we explain their structure:

- **TLWE ciphertexts:** The message m to be encrypted is encoded as an element \mathbb{T}_q . A mask $\mathbf{a} = (a_1, \dots, a_n)$ is drawn uniformly from \mathbb{T}_q^n and a noise error e is sampled from χ . Using the secret key $\mathbf{sk} = (s_1, \dots, s_n) \in \mathbb{B}^n$, the body of the ciphertext is defined by $b = \sum_{j=1}^n a_j \cdot s_j + m + e$. Finally, the TLWE ciphertext is $c = (\mathbf{a}, b)$. The decryption is performed by calculating the *phase*: $\phi(c) = b - \langle \mathbf{a}, \mathbf{s} \rangle = m + e$ and rounding to the closest center of sector.
- **TRLWE ciphertexts:** It has the same global structure as TLWE ones, except the mask \mathbf{a} is sampled from $\mathbb{T}_{N,q}[X]^k$, the secret key from $\mathbb{B}[X]^k$ and the error from $\mathbb{T}_{N,q}[X]$. In practice, it is common to choose $k = 1$.

During the bootstrapping phase presented in Section 2.5, another structure (the TRGSW ciphertext) is used but it will not be used in this paper. More details about TRGSW can be found in [20].

Two basics homomorphic operations are straightforward with these two structures: the component-wise sum of two TLWE (resp. TRLWE) ciphertexts c_1 and c_2 engenders a ciphertext c_3 encrypting the sum modulo p of the two underlying messages m_1 and m_2 . Moreover, the external product $\lambda \cdot c_1$ with $\lambda \in \mathbb{Z}$ also produces an encryption of the multiplication $[\lambda \cdot m_1]_p$.

In the framework introduced by this paper, only TLWE ciphertexts are used. The TRLWE ciphertexts are only required to understand the inner working of the programmable bootstrapping.

2.5 TFHE programmable bootstrapping (PBS)

As defined by Gentry in [18], the procedure of bootstrapping can be defined as the homomorphic evaluation of the decryption circuit. In the context of TFHE, the hardest part to compute is the rounding of the value to an element of \mathbb{T}_p by removing the noise. To achieve this homomorphically, it uses three procedures called **BlindRotate**, **SampleExtract** and **Functional Keyswitching**.

BlindRotate: The high level idea starts by homomorphically computing the phase $\mu \in \mathbb{Z}_q$ and reducing it to $\tilde{\mu} \in \mathbb{Z}_{2N}$ by computing $\tilde{\mu} = \left\lfloor \frac{\mu \cdot 2N}{q} \right\rfloor$. In practice N takes values between 2^{10} and 2^{13} so the most significant bits carrying the true value modulo p are preserved. Then, a polynomial called *accumulator* $v(X) \in \mathbb{Z}_{N,q}[X]$ is constructed. Note that the computation of $X^{-\tilde{\mu}} \cdot v(X)$ yields another polynomial of $\mathbb{Z}_{N,q}[X]$ of the form

$v_{\tilde{\mu}} + v_{\tilde{\mu}+1}X + \dots$. We select $v_j := \lfloor \frac{j\tilde{\mu}}{2^N} \rfloor$ so the blind rotation will output an encrypted version of the clean value of the message in the zero-degree coefficient. We do not explain here how this polynomial multiplication occurs, the reader is referred to [13] for a more elaborated explanation. The procedure outputs a TRLWE ciphertext of dimension k encrypting the polynomial $X^{-\tilde{\mu}} \cdot v(X)$. Note that the quotient polynomial of the ring has degree N but $\tilde{\mu}$ lives in \mathbb{Z}_{2N} so each coefficient of v_i can be reached with a multiplication by $X^{-\tilde{\mu}}$ and by $X^{\lfloor \frac{N-\tilde{\mu}}{2^N} \rfloor}$. In the latter case, the coefficient v_i gets negated because of the ring modulus $X^N + 1$: we'll refer to this problem as the *negacyclicity problem*. One way to prevent this issue is to ensure that the most significant bit of μ is fixed at 0 [20] but a recent work [14] proposes a more sophisticated way to solve this problem. In our case, we use a modified version of the accumulator detailed in Section 6.

SampleExtract: This step simply extracts the degree-zero coefficient of the previous polynomial. It takes as input the TRLWE ciphertext yielded by the **BlindRotate** step and outputs the TLWE ciphertext c' encrypting the original message m . However, this ciphertext is not immediately available for either further homomorphic computations or decryption, because it has a length $kN + 1$ instead of $n + 1$ (and as a consequence is encrypted under a different TLWE key).

Key switching: The previous step outputs the right value, but encrypted under a different set of parameters i.e. $c' \in \mathbb{Z}_q^{kN+1}$ while we are looking for $c \in \mathbb{Z}_q^n$. The only thing left is to convert c' to c , which requires *key switching keys* constructed from the secret key sk used at encryption. More details about this specific step can also be found in [13].

This “bland” procedure of bootstrapping simply refreshes the noise in the ciphertext to put it back at the “initial level”, but can be very simply turned into a *Programmable* bootstrapping. Specifically it can simultaneously evaluate homomorphically any function f on the input. To achieve this, at the construction of the accumulator, the coefficient v_j is replaced by their evaluation by the function $f(v_j)$. This feature is extremely powerful and is the core of the huge potential of TFHE.

2.6 Basics on Boolean Functions and Boolean Circuits

In this paper, we focus on the evaluation of Boolean functions with TFHE. A Boolean function has the form $f : \mathbb{B}^\ell \rightarrow \mathbb{B}$, with ℓ being called the *arity* of the function.

Definition 3. *The Algebraic Normal Form (ANF) of a Boolean function $f : \{0, 1\}^\ell \mapsto \{0, 1\}$ is a polynomial expression in which each term corresponds to a specific input combination of n variables. The ANF is defined as follows:*

$$f(x_1, x_2, \dots, x_\ell) = a_0 \oplus a_1 x_1 \oplus a_2 x_2 \oplus \dots \oplus a_{2^n - 1} x_1 x_2 \dots x_\ell$$

where: $a_0, a_1, a_2, \dots, a_{2^\ell - 1} \in \{0, 1\}$ are the Boolean coefficients and
 x_1, x_2, \dots, x_ℓ are called the Boolean variables

This result means that any Boolean function can be evaluated by the means of **AND** and **XOR** operations. In the following, we will focus on the implementation of Boolean circuits composed of these operations exclusively.

A Boolean function can be represented by its *truth table*, which is simply a table gathering all the possible inputs and the corresponding result of the application by the function. It can also be represented with a Boolean formula. A third representation is the *Boolean circuit*:

Definition 4. *A Boolean circuit associated to the Boolean function f is a finite Directed Acyclic Graph whose edges are wires and vertices are Boolean gates representing Boolean operations. We consider **AND***

gates and XOR gates, of fan-in 2 and fan-out 1. We also consider copy gates, of fan-in 1 and fan-out > 1 , that outputs several copies of its input. A circuit is further formally composed of input gates of fan-in 0 and fan-out 1, and output gates of fan-in 1 and fan-out 0.

Evaluating a ℓ -input m -output circuit consists in writing an input $\mathbf{x} \in \mathbb{B}^\ell$ in the input gates, processing the gates from input gates to output gates, then reading the outputs from the output gates.

This notion of Boolean circuit will be particularly useful in Section 5.

3 Boolean Encoding over \mathbb{Z}_p and Homomorphic Evaluation Strategy Between \mathbb{B} and \mathbb{Z}_p

To evaluate Boolean functions in TFHE, one could use the vanilla TFHE with $p = 2$. The problem is that the only evaluable function would be the XOR operation. To evaluate the other operators, the solution of [13] which is also implemented in the `tfhe-rs` library [27] that we studied is to take a larger p , specifically $p = 8$. This allows all the operations of the Boolean algebra to be carried out, however the negacyclicity problem introduced in Section 2.5 arises because 8 is even. Their solution to this issue is to keep a *bit of padding* fixed to zero, i.e. the values in \mathbb{Z}_p have their most significant bit fixed to zero. This restriction has a heavy impact on performances, because it requires a bootstrapping after each Boolean gate to make sure no data ever overflows in the most significant bit.

Our solution makes use of *odd* values for p , which allows us to remove this constraint of padding and to perform more operations without bootstrapping. To do so, we had to slightly adapt the bootstrapping procedure of TFHE to support odd moduli. We explain this tweak in Section 6.

Moreover, the PBS described in Section 2.5 takes only one input and so can only evaluate univariate functions. The common solution to evaluate multivariate functions is to concatenate several input ciphertexts into one by shifting the MSB of each input and to sum them all. The problem is that the number of message bits cannot grow too much because the other parameters of the LWE problem must grow accordingly, degrading the performances. As a consequence, the performances quickly degrades as the arity of the function increases. Our approach consists in removing the padding bit and using a combination of homomorphic additions before a PBS to evaluate a function for *any* number of inputs with the cost of a single PBS.

To this purpose, we propose a construction in which we embed Boolean values in \mathbb{Z}_p for p well-chosen values of p , forming an “intermediate homomorphic layer” between \mathbb{B} and \mathbb{Z}_q . In the following, we explain how we define such a layer, and we describe our new strategy to evaluate Boolean functions in a more efficient way without considering the circuit representation of the function.

3.1 Encoding of \mathbb{B} over \mathbb{Z}_p

To represent Boolean values over \mathbb{Z}_p , we use a mapping function that we call a *p-encoding*:

Definition 5. A *p-encoding* is a function $\mathcal{E} : \mathbb{B} \mapsto 2^{\mathbb{Z}_p}$ that maps the Boolean space to a subset of the discretized torus. A *p-encoding* is valid if and only if:

$$\begin{cases} \mathcal{E}(0) \cap \mathcal{E}(1) = \emptyset \text{ and} \\ \forall x \in \mathbb{Z}_p, \forall b \in \mathbb{B}, x \in \mathcal{E}(b) \iff [x + \frac{p}{2}]_p \notin \mathcal{E}(b) \text{ if } p \text{ is even.} \end{cases} \quad (1)$$

We call this last property relaxed negacyclicity.

In our approach when we need to encrypt a bit, we apply a *p-encoding* to embed it on \mathbb{Z}_p , then we encrypt the result using the classical setup of TFHE. When new values are freshly encrypted or produced by a PBS, only one element of \mathbb{Z}_p is chosen for each bit. We call such an encoding a canonical *p-encoding*:

Definition 6. A *p-encoding* \mathcal{E} is said canonical if and only if it is valid and $|\mathcal{E}(0)| = |\mathcal{E}(1)| = 1$




Fig. 1: Representation of two valid p -encodings. The green part represents $\mathcal{E}(1)$, and the red part $\mathcal{E}(0)$. Note that the relaxed negacyclity is respected by the p -encoding on the right-hand figure as p is even.

Let c be a ciphertext encoding a bit b under a p -encoding \mathcal{E} . Let us say that \mathcal{E} is an arbitrary valid encoding: its associated subsets can be any subset of \mathbb{Z}_p as long as the validity requirements of (1) are fulfilled. One can transform the ciphertext c into another ciphertext c' encoded under any **canonical** p -encoding (even under a different p) by simply performing a PBS.

Our goal is to represent the Boolean function we want to evaluate with a sum of p -encodings (we define what we mean by “sum of p -encoding” in the Section 3.2). When sums are carried out on ciphertexts (and so homomorphically on the underlying p -encodings), the sets $\mathcal{E}(0)$ and $\mathcal{E}(1)$ of the p -encodings may move, grow, shrink, but they should never overlap as it would result in a loss of information. As we removed the need of a bit of padding, we do not need to track a potential overflow of data (informally we say that ciphertexts are free to “go around the torus”). After the sum, the encoding can be reset to a canonical one with a PBS to allow further computation. Or, if the homomorphic computation is over, the result can be recovered by decrypting the ciphertext and checking in which partition the decrypted value lies.

The next subsection explains in further details the process of evaluating Boolean functions on \mathbb{Z}_p .

3.2 A New Strategy for Homomorphic Boolean Evaluation

In the following, we consider two Boolean variables x and y and their two respective encodings over \mathbb{Z}_p :

$$\mathcal{E}_x = \begin{cases} 0 \mapsto \{\alpha_i\}_{0 \leq i \leq l_0} \\ 1 \mapsto \{\beta_i\}_{0 \leq i \leq l_1} \end{cases} \quad \text{and} \quad \mathcal{E}_y = \begin{cases} 0 \mapsto \{\alpha'_i\}_{0 \leq i \leq l'_0} \\ 1 \mapsto \{\beta'_i\}_{0 \leq i \leq l'_1} \end{cases} \quad (2)$$

Let f be a bivariate Boolean function and let us construct two sets \mathcal{P}_0 and \mathcal{P}_1 such that:

$$\mathcal{P}_b = \{[\gamma + \delta]_p \mid (\gamma, \delta) \in \mathcal{E}_x(b_x) \times \mathcal{E}_y(b_y) \text{ and } f(b_x, b_y) = b\} \quad \forall b \in \mathbb{B}. \quad (3)$$

We say that the sum of p -encodings $\mathcal{E}_x + \mathcal{E}_y$ is *suitable for the evaluation of f* if and only if $\mathcal{P}_0 \cap \mathcal{P}_1 = \emptyset$. The definition can be generalized to any number of arguments ℓ for f . Of course, for a given f , not all the pairs of encodings may be suitable. And finding such correct encodings might be a complicated process.

If \mathcal{E}_x and \mathcal{E}_y are suitable for f , then one can use the computed sets \mathcal{P}_b to construct a new p -encoding

$$\mathcal{E}_z = \begin{cases} 0 \mapsto \mathcal{P}_0 \\ 1 \mapsto \mathcal{P}_1 \end{cases}$$

that encodes the bit $f(x, y)$. As \mathcal{E}_z is valid, then the clear value of the bit can be recovered by decryption, or further computations can be performed without the need of a bootstrapping.

Definition 7. Let $f : \mathbb{B}^\ell \mapsto \mathbb{B}$ be a Boolean function and let $\mathcal{E} = (\mathcal{E}_1, \dots, \mathcal{E}_\ell)$ be a vector of p -encodings. We define $f(\mathcal{E})$ by:

$$f(\mathcal{E}) = \begin{cases} 0 \mapsto \mathcal{P}_0 \\ 1 \mapsto \mathcal{P}_1 \end{cases}$$

with:

$$\mathcal{P}_b = \left\{ \left[\sum_{i=1}^l \gamma_i \right]_p \mid (\gamma_1, \dots, \gamma_l) \in \prod_{i=1}^l \mathcal{E}_i(b_i) \text{ and } f(b_1, \dots, b_l) = b \right\} \forall b \in \mathbb{B}$$

Let us stress that $f(\mathcal{E})$ is a valid p -encoding if and only if $\mathcal{P}_0 \cap \mathcal{P}_1 = \emptyset$.

In the following, we illustrate the latter definition with the two Boolean operators $\&$ and \oplus . The p -encoding resulting of the function $f : (x, y) \mapsto x \& y$ is:

$$\mathcal{E}_{\&} = \begin{cases} 0 \mapsto \{\alpha_i + \alpha'_j\}_{\substack{0 \leq i \leq l_0 \\ 0 \leq j \leq l'_0}} \cup \{\alpha_i + \beta'_j\}_{\substack{0 \leq i \leq l_0 \\ 0 \leq j \leq l'_1}} \cup \{\alpha'_i + \beta_j\}_{\substack{0 \leq i \leq l'_0 \\ 0 \leq j \leq l_1}} \\ 1 \mapsto \{\beta_i + \beta'_j\}_{\substack{0 \leq i \leq l_1 \\ 0 \leq j \leq l'_1}} \end{cases} \quad (4)$$

and the p -encoding resulting of the operation $f : (x, y) \mapsto x \oplus y$ is:

$$\mathcal{E}_{\oplus} = \begin{cases} 0 \mapsto \{\alpha_i + \alpha'_j\}_{\substack{0 \leq i \leq l_0 \\ 0 \leq j \leq l'_0}} \cup \{\beta_i + \beta'_j\}_{\substack{0 \leq i \leq l_0 \\ 0 \leq j \leq l'_1}} \\ 1 \mapsto \{\alpha_i + \beta'_j\}_{\substack{0 \leq i \leq l_0 \\ 0 \leq j \leq l'_1}} \cup \{\alpha'_i + \beta_j\}_{\substack{0 \leq i \leq l'_0 \\ 0 \leq j \leq l_1}} \end{cases} \quad (5)$$

Figure 2 further illustrates this construction for these two operations.

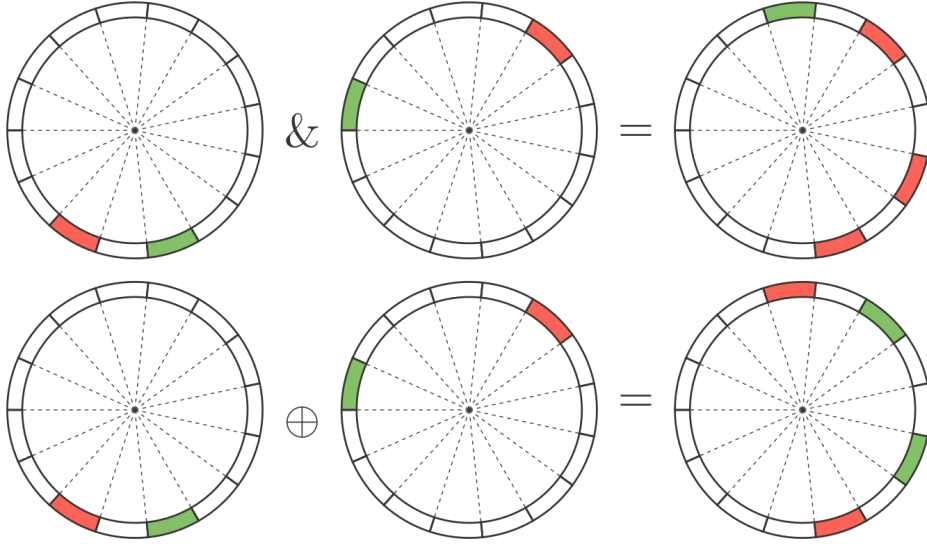


Fig. 2: Starting from two canonical encodings, we produce two new p -encodings corresponding to the results of the $\&$ and the \oplus operations.

To wrap up, here is our proposed framework to evaluate a Boolean function $f : \mathbb{B}^\ell \mapsto \mathbb{B}$ given a vector of suitable p -encodings $\mathcal{E} = (\mathcal{E}_1, \dots, \mathcal{E}_l)$:

1. Encrypt each input b_i with some canonical p -encoding \mathcal{E}_i into a ciphertext c_i such that $\mathcal{E}_{sum} = f(\mathcal{E}_1, \dots, \mathcal{E}_l)$ is a valid encoding.
2. For a Boolean function f to be evaluated on b_1, \dots, b_l , compute homomorphically the sum of the ciphertexts $c = c_1 + \dots + c_l$. This yields an encryption of $b = f(b_1, \dots, b_l)$, encoded with a valid p -encoding $\mathcal{E}_{sum} = f(\mathcal{E}_1, \dots, \mathcal{E}_l)$.

3. (a) If the result is directly required by the client, c is returned as ciphertext which can be decrypted to get the result in \mathbb{Z}_p and associated to the right Boolean value.
- (b) If the result is reused in further homomorphic computations, a PBS calculating $\mathcal{E}_{new} \circ \mathcal{E}_{sum}^{-1}$ on the result is computed, with \mathcal{E}_{new} a new canonical p -encoding. The resulting value can then be used as an input for a next Boolean function.

Of course, for a given p and a given f , a vector of p -encodings $(\mathcal{E}_1, \dots, \mathcal{E}_l)$ for which $\mathcal{E}_{sum} = f(\mathcal{E}_1, \dots, \mathcal{E}_l)$ is valid may not exist. Two solutions can be considered:

- Increasing the value of p (e.g $p \geq 2^\ell$ always works).
- “Splitting” the function into a graph of subfunctions, evaluating each subfunction with the above method and using a PBS per subfunction to reset its output p -encodings.

The question of finding suitable p -encodings for a given f is treated in Section 4. The question of efficiently splitting a function is treated in Section 5.

Example: We illustrate our approach with a simple working example: let f be a basic multiplexing function, such that

$$f(a, b, c) = \begin{cases} a & \text{if } c = 1 \\ b & \text{if } c = 0 \end{cases}$$

Instead of leveraging its Boolean representation $f(a, b, c) = a \& c \oplus b \& \bar{c}$, which would cost 3 PBS with the approach of [13], our strategy consists in:

1. Encrypting the bits with the 7-encodings:

$$\mathcal{E}_a = \mathcal{E}_b = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{1\} \end{cases} \quad \text{and} \quad \mathcal{E}_c = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{2\} \end{cases}$$

2. Applying the function f on the 7-encodings by summing the ciphertexts, producing a valid 7-encoding:

$$\mathcal{E}_{sum} = \begin{cases} 0 \mapsto \{0, 1, 2, 5\} \\ 1 \mapsto \{3, 4, 6\} \end{cases}$$

At this point, only sums have been performed on the ciphertexts.

3. With one PBS, resetting the result to a target canonical p -encoding (with any p), for example

$$\mathcal{E}_{new} = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{1\} \end{cases} \quad \text{with } p = 7$$

A visualization of this procedure can be found in Figure 3.

We can keep doing that with subsequent Boolean functions. Alternatively we can return the ciphertext under \mathcal{E}_{sum} without PBS which can then be decrypted and decoded as the result of the function f .

3.3 Encoding Switching

One can homomorphically switch the encoding of a ciphertext. Let us begin with some trivial cases:

Property 8. Let x be a ciphertext encoded under $\mathcal{E}_x = \begin{cases} 0 \mapsto \{\alpha_i\}_{0 \leq i \leq l_0} \\ 1 \mapsto \{\beta_i\}_{0 \leq i \leq l_1} \end{cases}$ and $a \in \mathbb{Z}_p$ a constant. The encoding of x can be switched to:

$$\mathcal{E}'_x = \begin{cases} 0 \mapsto \{[\alpha_i + a]_p\}_{0 \leq i \leq l_0} \\ 1 \mapsto \{[\beta_i + a]_p\}_{0 \leq i \leq l_1} \end{cases}$$

by an homomorphic addition of the ciphertext x and the clear value a .

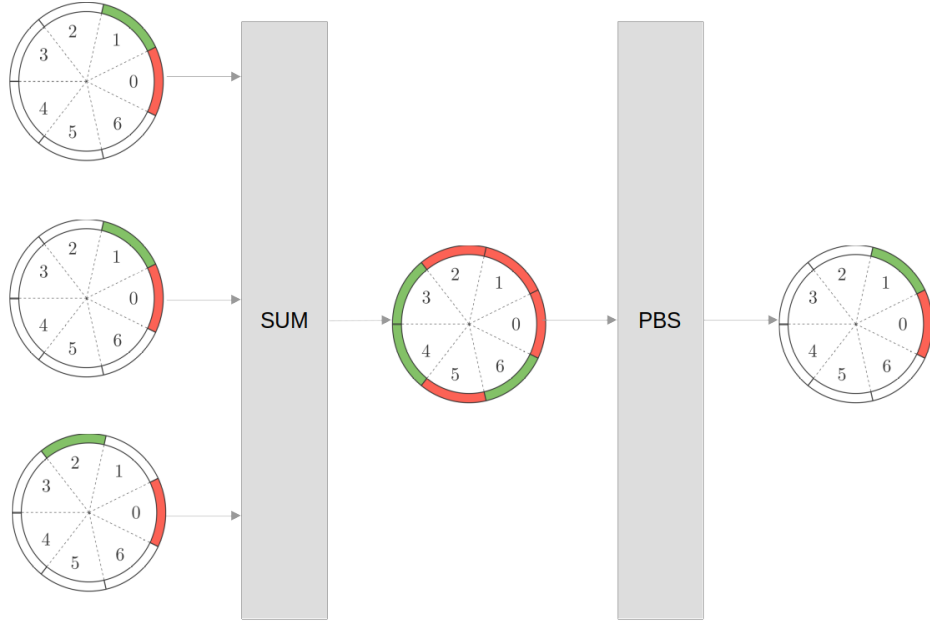


Fig. 3: Illustration of an execution of the framework for the multiplexing function.

Proof. All the elements of $\mathcal{E}'_x(0)$ (resp. $\mathcal{E}'_x(1)$) are offset by exactly a from their counterparts in $\mathcal{E}_x(0)$ (resp. $\mathcal{E}_x(1)$). Thus, if the original encoding \mathcal{E}_x was valid, then $\mathcal{E}_x(0) \cap \mathcal{E}_x(1) = \emptyset$. So we trivially get $\mathcal{E}'_x(0) \cap \mathcal{E}'_x(1) = \emptyset$ and thus the validity of \mathcal{E}'_x . \square

Property 9. Let x be a ciphertext encoded under the p -encoding: $\mathcal{E}_x = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{1\} \end{cases}$ and let $a \in \mathbb{Z}_p^*$. Then, it

can be switched to: $\mathcal{E}'_x = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{a\} \end{cases}$ by a simple homomorphic multiplication between the ciphertext x and the clear constant a . Note that the same holds for the same p -encodings but with $\mathcal{E}(0)$ and $\mathcal{E}(1)$ swapped.

Proof. The property is trivial by the linear homomorphism of the TFHE scheme. \square

Other operations of encoding switching can be carried out with a homomorphic multiplication by a plaintext, but their validity depends of the value of p . Lemma 12 in Section 4 addresses the case of prime values of p .

Finally, a programmable bootstrapping can also be seen as an encoding switching. It turns any valid encoding into another valid encoding, even with a different modulus p . A good example of that is the step 3.(b) of the protocol described in Section 3.2. Let us explain this analogy:

Recall that the programmable bootstrapping of TFHE is a homomorphic evaluation of look-up table: for a message $m \in \mathbb{Z}_p$ encrypted under a ciphertext c and a function f , the PBS outputs an encryption c' of $f(m)$. Let us re-use the p -encodings \mathcal{E}_{sum} and \mathcal{E}_{new} introduced in Section 3.2. The abstract look-up table has p lines, indexed from 0 to $p - 1$. The lines whose indexes belong to $\mathcal{E}_{sum}(0)$ (resp. $\mathcal{E}_{sum}(1)$) are filled with an element of $\mathcal{E}_{new}(0)$ (resp. $\mathcal{E}_{new}(1)$). Other lines are never reached and can be filled with an arbitrary placeholder value. This way, evaluating the look-up table turns an encryption of the message under \mathcal{E}_{sum} into an encryption under \mathcal{E}_{new} . See Sections 2.5 and 6.2 for a more in-depth insight on the actual procedure of programmable bootstrapping.

4 Finding Valid Sets of Input Encodings

Let $f : \mathbb{B}^\ell \mapsto \mathbb{B}$ a Boolean function with ℓ entries. This section addresses the problem of finding, for a given p , a set of ℓ p -encodings allowing the homomorphic evaluation of f with one single bootstrapping.

4.1 Reduction of the Search Space

While exhaustive search is a first option, it quickly becomes impractical due to the explosion of the number of possibilities as p grows. As a consequence, a reduction of the search space is needed without leaving out a potential solution.

We introduce two lemmas that will be used to reduce the search space:

Lemma 10. *Let $f : \mathbb{B}^\ell \mapsto \mathbb{B}$ and let $(\mathcal{E}_1, \dots, \mathcal{E}_\ell)$ a vector of p -encodings suitable for f and having the form:*

$$\forall i \in \{1, \dots, \ell\}, \mathcal{E}_i = \begin{cases} 0 \mapsto \{x_j^{(i)}\}_{1 \leq j \leq l_0^{(i)}} \\ 1 \mapsto \{y_j^{(i)}\}_{1 \leq j \leq l_1^{(i)}} \end{cases} . \text{ Then any vector of canonical } p\text{-encodings } (\mathcal{E}'_1, \dots, \mathcal{E}'_\ell) \text{ of the}$$

form: $\forall i \in \{1, \dots, \ell\}, \mathcal{E}'_i = \begin{cases} 0 \mapsto \{x^{(i)}\} \\ 1 \mapsto \{y^{(i)}\} \end{cases}$ with $x^{(i)} \in \mathcal{E}_i(0)$ and $y^{(i)} \in \mathcal{E}_i(1)$ is suitable for the function f as well.

Proof. Let us assume that the vector $\mathcal{E} = (\mathcal{E}_1, \dots, \mathcal{E}_\ell)$ of Lemma 10 is suitable for the function f . Then the sets \mathcal{P}_0 and \mathcal{P}_1 constructed like in Equation 3 are disjoint. Now, let us consider the vector of canonical p -encodings $\mathcal{E}' = (\mathcal{E}'_1, \dots, \mathcal{E}'_\ell)$ respecting the property:

$$\forall b \in \mathbb{B}, \forall i \in \{0, \dots, \ell\}, \mathcal{E}'_i(b) \subset \mathcal{E}_i(b).$$

As a consequence, if we build the sets \mathcal{P}'_0 and \mathcal{P}'_1 relative to the encodings \mathcal{E}' , then we naturally get $\mathcal{P}'_0 \subset \mathcal{P}_0$ and $\mathcal{P}'_1 \subset \mathcal{P}_1$. So we get $\mathcal{P}'_0 \cap \mathcal{P}'_1 = \emptyset$, proving Lemma 10. \square

Lemma 11. *Let $f : \mathbb{B}^\ell \mapsto \mathbb{B}$ and let $(\mathcal{E}_1, \dots, \mathcal{E}_\ell)$ a vector of p -encodings suitable for f and of the form:*

$$\forall i \in \{1, \dots, \ell\}, \mathcal{E}_i = \begin{cases} 0 \mapsto \{x^{(i)}\} \\ 1 \mapsto \{y^{(i)}\} \end{cases} \text{ Then any vector of canonical } p\text{-encodings } (\mathcal{E}'_1, \dots, \mathcal{E}'_\ell) \text{ of the form: } \forall i \in$$

$\{1, \dots, \ell\}, \mathcal{E}'_i = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{y^{(i)} - x^{(i)}\} \end{cases}$ *is suitable for the function f as well.*

Proof. Let $f : \mathbb{B}^\ell \mapsto \mathbb{B}$ be a function and \mathcal{E} be a vector of canonical p -encodings $(\mathcal{E}_1, \dots, \mathcal{E}_\ell)$ suitable for f with:

$$\forall i \in \{1, \dots, \ell\}, \mathcal{E}_i = \begin{cases} 0 \mapsto \{x^{(i)}\} \\ 1 \mapsto \{y^{(i)}\} \end{cases} .$$

Let us build the sets \mathcal{P}_0 and \mathcal{P}_1 according to Equation 3. Each element of these sets is the sum of exactly one element of each p -encoding, that is to say an element $\mathcal{E}_i(0) \cup \mathcal{E}_i(1)$.

Let us pick an indice $k \in \{1, \dots, \ell\}$, a value $a \in \mathbb{Z}_p$ and replace \mathcal{E}_k in the vector \mathcal{E} by:

$$\mathcal{E}'_k = \begin{cases} 0 \mapsto \{x^{(i)} - a\} \\ 1 \mapsto \{y^{(i)} - a\} \end{cases}$$

By using the Property 8, we directly have $\mathcal{P}'_0 \cap \mathcal{P}'_1 = \emptyset$ from $\mathcal{P}_0 \cap \mathcal{P}_1 = \emptyset$ (by suitability of the encodings for f).

By iterating this procedure on each of the ℓ elements of \mathcal{E} , and by picking each time $a = -x^{(i)}$, we prove Lemma 11. \square

Using both Lemmas 10 and 11, we can restrict the search to the encodings of the form

$$\mathcal{E}_i = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{d_i\} \end{cases}$$

with $d_i \neq 0$ without any loss of generality.

Moreover, we restrict the solution further: we only consider p -encodings with p odd and prime. The choice of an odd p allows to free ourselves from the negacyclicity constraint. To explain the constraint of primality, we introduce the following lemma, that allows to drastically improve the performances of the search:

Lemma 12. *Let p be a prime and $f : \mathbb{B} \rightarrow \mathbb{B}$ be a Boolean function and let $\mathcal{E} = (\mathcal{E}_1, \dots, \mathcal{E}_l)$ be p -encodings suitable for f with: $\forall i \in \{1, \dots, l\}, \mathcal{E}_i = \begin{cases} 0 \mapsto \{x^{(i)}\} \\ 1 \mapsto \{y^{(i)}\} \end{cases}$. For every $a \in \mathbb{Z}_p^*$, the vector of p -encodings $\mathcal{E}' = (\mathcal{E}'_1, \dots, \mathcal{E}'_l)$ with: $\mathcal{E}'_i = \begin{cases} 0 \mapsto \{[a \cdot x^{(i)}]_p\} \\ 1 \mapsto \{[a \cdot y^{(i)}]_p\} \end{cases}$ is suitable for f as well.*

Proof. As p is prime, the multiplication by a is a bijection from \mathbb{Z}_p to \mathbb{Z}_p . If we construct the two sets \mathcal{P}_0 and \mathcal{P}_1 of Equation (3) with the vector \mathcal{E} , then we have $\mathcal{P}_0 \cap \mathcal{P}_1 = \emptyset$. By applying a multiplication by a fixed a in \mathbb{Z}_p^* to each of the elements $x^{(i)}$ and $y^{(i)}$, the sets stay disjoint (because of the bijection) so the new set of encodings \mathcal{E}' is suitable for f as well. \square

As a consequence, if p is prime (which we shall always choose in practice), any solution can be turned into a solution with $d_1 = 1$ by simply multiplying all the p -encodings of the solution by $[d_1^{-1}]_p$. So we can fix $d_1 = 1$ without any loss of generality, reducing drastically the size of the search space.

Moreover, one may think that the dimensioning of the sets of TFHE parameters is closely related to p . In practice, we actually select those depending only on the bit-size of p . Thus, the restriction to prime modulus has no impact on performances, as one can find prime numbers of any bit size.

4.2 Formalization of the Search Problem

According to the lemmas from Section 4.1, we can reduce the problem of finding a vector of p -encodings $(\mathcal{E}_1, \dots, \mathcal{E}_l)$ such that $f(\mathcal{E}_1, \dots, \mathcal{E}_l)$ is valid to the problem of finding a vector $\mathbf{d} = (d_1, \dots, d_l)$ such that

$$\mathcal{E}_i = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{d_i\} \end{cases} \quad \text{and } f(\mathcal{E}_1, \dots, \mathcal{E}_l) \text{ is valid. In the following, we describe an algorithm to find such a vector } \mathbf{d}.$$

We denote V the matrix of elements of \mathbb{B} of shape $2^l \times l$ gathering all the possible sequences of entries for the function f :

$$V = \begin{pmatrix} 0 & \dots & 0 & 0 \\ 0 & \dots & 0 & 1 \\ \vdots & \ddots & \vdots & \vdots \\ 1 & \dots & 1 & 1 \end{pmatrix}$$

Also, we denote by \mathbf{b} the vector of all the outputs of the function f , sorted in same order as the rows of V . Thus, we have: $\forall i \in \{1, \dots, 2^l\}, b_i = f(V_i)$ for V_i the i th row of V . Let us define the vector \mathbf{r} as: $\mathbf{r} = V\mathbf{d}$. To make \mathbf{d} a solution of the problem, \mathbf{r} has to verify the following property:

$$\forall i, j \in \{1, \dots, 2^l\}, f(V_i) \neq f(V_j) \implies r_i \neq r_j \quad (6)$$

An alternative formulation is: we look for two disjoint subsets \mathcal{P}_0 and \mathcal{P}_1 of \mathbb{Z}_p , such that: $f(V_i) = b \iff r_i \in \mathcal{P}_b$.

The following section describes an algorithm finding a solution to this problem.

4.3 Algorithm

We start by constructing two sets \mathcal{F} and \mathcal{T} such that: $\mathcal{F} = \{V_i | b_i = 0\}$ and $\mathcal{T} = \{V_i | b_i = 1\}$. Each line V_i represents a linear combination of the d_j 's, that verifies: $r_i = \sum_{j=0}^l V_{ij} \cdot d_j \pmod p$. The values r_i produced by the elements of \mathcal{F} must be different from the ones produced by \mathcal{T} . As a consequence, we can write: $\forall (V_i, V_j) \in \mathcal{F} \times \mathcal{T}, \sum_{k=0}^l V_{ik} \cdot d_k \neq \sum_{k=0}^l V_{jk} \cdot d_k$, which is equivalent to write: $\forall (V_i, V_j) \in \mathcal{F} \times \mathcal{T}, \sum_{k=0}^l (V_{ik} - V_{jk}) \cdot d_k \neq 0$. So we can rewrite our constraints in the set $\mathcal{C} = \{V_i - V_j | (V_i, V_j) \in \mathcal{F} \times \mathcal{T}\}$. \mathcal{C} contains vectors with coordinates in $\{0, 1, -1\}$ representing linear combinations that have to be non-zero. Note that if an element of the set \mathcal{C} is the opposite of an other, it does not bring further constraint and can thus be safely discarded from the set.

The use of a set in the implementation at this point of the algorithm allows to remove a lot of duplicate constraints and to simplify the next step. Then, the problem reduces to solving a “linear system of inequalities” in the ring \mathbb{Z}_p :

$$\begin{cases} c_1^{(1)} \cdot d_1 + \dots + c_l^{(1)} \cdot d_l \neq 0 \pmod p \\ c_1^{(2)} \cdot d_1 + \dots + c_l^{(2)} \cdot d_l \neq 0 \pmod p \\ \vdots \end{cases} \quad \text{with } c_i^{(j)} \in \{0, \pm 1\}$$

After filtering, we pack all the elements of \mathcal{C} in ℓ matrices $\{C_i\}_{1 \leq i \leq \ell}$ (each row being a linear combination), where the matrix C_i packs all the constraints involving only the i first inputs (i.e. all the coefficients of column index greater than i are zeros).

We then perform a recursive search (Algorithm 1), affecting at each step of depth i a possible value d_i for the i -th input. To do so, we call Algorithm 2 to construct the set of all possible values complying with the constraints of the matrix C_i and the previously set values for the preceding inputs. If we reach a dead-end, we backtrack by deleting the preceding input and assigning it the next possible value. Algorithms 1 and 2 formalize this idea: Algorithm 1 is a basic recursive backtracking algorithm using calls to the set construction function (Algorithm 2) to get the possibilities for the next value of \mathbf{d} . The latter, when called at depth $j + 1$, takes as input the j values already computed at higher depth for \mathbf{d} and the matrix of constraints C_{j+1} . Each line of C_{j+1} creates a (potentially duplicate) forbidden value for d_{j+1} , these values are all computed and the complement of this set in \mathbb{Z}_p is returned by the algorithm (i.e. the set for possible values for d_{j+1} at this point of the search).

Theorem 13. *Running Algorithm 1 with increasing values of p ensures that the first solution \mathbf{d} found is optimal for the function f , i.e. the solution works and its associated p is the smallest as possible.*

Optimizations: Several optimizations are possible to improve the performances of the search. First, in Algorithm 2, one can check the size of the set \bar{S} at each iteration and stop as soon as the size of the set is p . Such a set means that a dead-end has been reached and that no value will be returned by the function. Then, one can leverage symmetries existing in the table but also in the function. For example, if we consider the function $f : (x, y) \rightarrow x \oplus y$, the two variables x and y have symmetric roles. Thus, if the pair of encodings $(\mathcal{E}_x, \mathcal{E}_y)$ is valid, then the pair $(\mathcal{E}_y, \mathcal{E}_x)$ is valid as well. As a consequence, one can arbitrarily set $d_x \leq d_y$ and removing half the possibilities for (x, y) .

Development of an heuristic: This algorithm of the previous section is deterministic and finds any existing set of encodings compliant with the function f . Nevertheless, it is not suitable for big functions for several reasons. First, the complexity of the algorithm is directly correlated with the number ℓ of inputs, and for big functions, higher values for p are required, which makes the growth of computation even worse. Moreover, the right value for p is not known *a priori*, so we have to run the full algorithm for each value of p until we find one that works. For these reasons, we might prefer an efficient heuristic over the previous algorithm in some contexts. In Section 4.4, we define such a heuristic which allows to drastically improve the performance by executing directly the algorithm with realistic values for p .

Algorithm 1 Recursive function `search` that adds an element to the vector \mathbf{d}

Require:

$\mathbf{d} := (d_i)_{1 \leq i \leq j}$ \triangleright The vector of values for the inputs already computed
 $\{C_i \mid i \in \{0, \dots, \ell - 1\}\}$ \triangleright The matrices of constraints, pre-computed
 $p \in \mathbb{N}^*$ \triangleright the modulus of the input encodings
 $\ell \in \mathbb{N}^*$ \triangleright The target number of encodings required

Ensure: f is evaluable using the encodings \mathbf{d} .

```
if  $j = \ell$  then
  return  $\mathbf{d}$   $\triangleright$  Base case of recursion when a complete solution has been found
else
   $\mathcal{P} \leftarrow \text{get\_possible\_values}(\mathbf{d}, C_{j+1}, p)$   $\triangleright$  Retrieving the set of possible values for  $d_k$ .

  for  $x \in \mathcal{P}$  do
     $\mathbf{d} \leftarrow (\mathbf{d} \parallel x)$   $\triangleright$  Affecting one of the possible value to  $d_k$ 
     $\mathbf{d}_{sol} \leftarrow \text{search}(\mathbf{d}, C, p, \ell)$   $\triangleright$  Recursively calling the algorithm
    if  $\mathbf{d}_{sol} \neq \perp$  then
      return  $\mathbf{d}_{sol}$   $\triangleright$  If a final solution has been found, propagating it to the higher level
    else
       $\mathbf{d} \leftarrow \mathbf{d}[:j+1]$   $\triangleright$  If the previous call failed, we remove the last value and try an other one.
  end if
end for
return  $\perp$   $\triangleright$  If all of the possibilities have been tested and none of them work, we need to backtrack
```

Algorithm 2 Function `get_possible_values` that builds the set of possible values for the next slot of \mathbf{d} given the slots already filled in.

Require:

$\mathbf{d} := (d_i)_{\{1 \leq i \leq j\}}$ \triangleright The set of values for the inputs already computed. Note that d_1 is fixed to 1
 C_{j+1} \triangleright The matrix of constraints of this step, pre-computed
 $p \in \mathbb{N}^*$ \triangleright the modulus of input encoding

Ensure: The set S contains only values suitable for the $j + 1$ -th slot of \mathbf{d} .

```
 $\bar{S} \leftarrow \{\}$   $\triangleright \bar{S}$  is the set of forbidden values for  $d_{j+1}$ 
for  $c \in C_k$  do
   $\bar{c} \leftarrow c[j+1]$   $\triangleright$  We retrieve the  $(j+1)$ th coefficient of the inequation  $c$ 
   $\bar{S} \leftarrow \bar{S} \cup \{-\bar{c} \cdot \sum_{k=0}^j c_k \cdot d_k\}$   $\triangleright$  We compute the value forbidden by  $c$ 
end for
 $S \leftarrow \mathbb{Z}_p \setminus \bar{S}$ 
return  $S$ 
```

4.4 An Efficient Sieving Heuristic to Find Suitable Encodings

The algorithm of the previous section is deterministic and finds any existing set of encodings compliant with the function f . Nevertheless, it is not suitable for big functions for several reasons. First, the complexity of the algorithm is directly correlated with the number ℓ of inputs, and for big functions, higher values for p are required, which makes the growth of computation even worse. Moreover, the right value for p is not known *a priori*, so we have to run the full algorithm for each value of p until we find one that works. For these reasons, we might prefer an efficient heuristic over the previous algorithm in some contexts. In the following, we define such a heuristic which allows to drastically improve the performance by executing directly the algorithm with realistic values for p .

Let us consider a function $f : \mathbb{B}^\ell \mapsto \mathbb{B}$ of matrix of constraints $C = (C_j^{(i)})_{\substack{1 \leq i \leq n_j \\ 1 \leq j \leq \ell}}$ and its associated system of linear inequalities:

$$\begin{cases} c_1^{(1)} \times d_1 + c_2^{(1)} \times d_2 + \dots + c_\ell^{(1)} \times d_\ell \neq 0 \pmod p \\ c_1^{(2)} \times d_1 + c_2^{(2)} \times d_2 + \dots + c_\ell^{(2)} \times d_\ell \neq 0 \pmod p \\ \dots \end{cases}$$

The principle is to sample random values in \mathbb{Z} (with some large bound) and affect them to the d_j 's. If all the corresponding values for all the $C_i = \sum_{j=1}^{\ell} c_j^{(i)} \times d_j$ are not divisible by a value p , then the vector $(d_j \pmod p \mid j \in \{1, \dots, \ell\})$ is a solution of the system of inequalities generated by C .

To reduce the amount of samples required to find a solution, we want to avoid sampling trivially wrong sets of d_j 's. For example, if all the d_j 's are themselves divisible by p , then the C_i 's will all be divisible as well. To tackle this problem, we perform the sampling across *prime numbers in \mathbb{Z}* .

Algorithm 3 Sample a solution \mathbf{d} in \mathbb{Z} for a function f and returns a possible value for p .

Require:

- $\{C_i\}_{1 \leq i \leq n}$ ▷ The lines of the matrix of constraints C of the function f
- P ▷ The sets of possible values for p to be tested
- D ▷ The sets of possible values in \mathbb{Z} to assign to the d_i 's. All these elements are big primes

Ensure: f is possible to evaluate using a modulus greater or equal than p .


- $\mathbf{d} \xleftarrow{\$} D$ ▷ Sample random prime values in \mathbb{Z} and assign it to $\mathbf{d} = (d_1, \dots, d_\ell)$
 - $\mathbf{r} = C \times \mathbf{d}$ ▷ \mathbf{r} is the right member of the system
 - for** $p \in P$ **do**
 - if** $0 \in [\mathbf{r}]_p$ **then** ▷ If p divides one of the coordinates of \mathbf{r}
 - $P \leftarrow P \setminus \{p\}$ ▷ This value of p is incorrect
 - end if**
 - end for**
 - if** $|P| > 0$ **then**
 - return** $\min(P)$ ▷ Returns the smallest possible value for p , if any.
 - end if**
-

Running this algorithm several times and keeping the smallest returned value for p , one gets an upper bound on the minimum p required to evaluate a function with our framework. Note that, on the contrary of the deterministic search algorithm, this heuristic does not require a prime p .

Example: Let us consider the s-box of the block cipher ASCON. We study this s-box in more details and provide an exact optimized solution for its homomorphic evaluation in Section 7.4. Here, we apply Algorithm 3 on the five functions generating the five output bits and monitor the results until we gather $N = 10000$ non-zero possible values for p .

The figure 4a shows the repartition of the returned values of p by the algorithm during these N runs on the first subfunction. The optimal value of p found by the deterministic approach of Section 4.3 is 17 so the

upper bound 19 is pretty close, despite being rarely found by the algorithm. Also, the figure 4b shows 21 (the second best solution found by the sieving) is almost instantly found by the algorithm.



(a) The outputs of 10000 runs of the Algorithm 3 for the first subfunction of the Ascon s-box

(b) Number of iterations required to get a solution for a given value of p

In the process of finding the smallest p possible and a correct vector of p -encoding to evaluate a function f , this heuristic is really efficient to get a tight upper bound on the value of p .

5 Scaling our Approach to any Boolean Circuit

The new method we introduced gives good results for simple functions but suffers the following limitations:

1. For a Boolean function with a high number of inputs, the search algorithm may be very time-consuming.
2. Some functions simply do not have any solution for acceptable values for p ($p < 32$ for example) and thus are not evaluable in one single bootstrapping.

As a consequence, we need a solution to extend our framework to these cases. In this section, we propose a strategy to leverage the circuit representation of a “tough” function f to find a strategy of homomorphic evaluation with as few bootstrappings as possible.

5.1 Graph of Subcircuits

Let $f : \mathbb{B}^\ell \rightarrow \mathbb{B}$ be a Boolean function, and let \mathcal{F} be a Boolean circuit representing f . A few preliminaries about Boolean circuits can be found in Section 2.6. Let us describe the layout of the circuit \mathcal{F} . It has l input wires, denoted by $\{y_j\}_{1 \leq j \leq \ell}$, and the output wire is denoted by z . The intermediary wires are denoted by $\{t_j\}_{1 \leq j \leq \theta}$. The Boolean operation gates are of fan-out 1. We introduce a *copy gate* which takes one wire as input and simply outputs $k > 1$ copies of the input.

Our goal is to split the circuit into a directed acyclic graph \mathcal{G} , whose vertexes are subcircuits $\{\mathcal{F}_1, \dots, \mathcal{F}_k\}$ and whose edges connect the outputs of a subcircuit with the input of another. Each subcircuit \mathcal{F}_i represents a subfunction $f_i : \mathbb{B}^{l_i} \mapsto \mathbb{B}$ that is evaluable in one single bootstrapping with our framework.

We use the same notations to refer to the elements of a subcircuit \mathcal{F}_i and we index them with i . The output of \mathcal{F}_i is denoted by $z^{(i)}$ and its inputs by $\{y_j^{(i)}\}_{1 \leq j \leq \ell}$ and so on.

The graph is valid for f with respect to modulus p if the following properties are satisfied:

- Each subcircuit \mathcal{F}_i has only one output $z^{(i)}$.
- For a subcircuit \mathcal{F}_i , all its inputs are either inputs of the whole circuit or outputs of other subcircuits of the graph. We can write this property as:

$$\{y_j^{(i)}\}_{1 \leq j \leq l^{(i)}} \subset \left(\{y_j\}_{1 \leq j \leq \ell} \cup \{z^{(j)}\}_{1 \leq j < i} \right)$$

Thus, the indexing of the \mathcal{F}_i 's respects the topological order of the graph, i.e. no gates of \mathcal{F}_i has a child in any of the \mathcal{F}_j , with $j < i$.

- All the Boolean functions f_i represented by the subcircuits \mathcal{F}_i are evaluable in a single bootstrapping with modulus p with our proposed method.
- The last subcircuit \mathcal{F}_c of the graph has z (the output of the main circuit) for output: $z^{(c)} = z$.

To homomorphically evaluate the function f , we evaluate each subcircuits with one bootstrapping for each of them and get the final result. In order to reduce the cost of evaluation for a given p , the goal is hence to find the *smallest* valid graph possible in terms of number of subcircuits. Taking a greater value of p produces a different graph that may be shorter (as subcircuits might be larger), but the timings of bootstrapping in this graph might on the other hand be greater. One can therefore run the search for different values of p and keep the most efficient setup among the possible graphs (which might depend on the target platform).

5.2 Heuristics to Find a Small Graph

Finding such a graph can be done by exhaustively evaluating all the possible subcircuits with our method introduced in Section 4, and then find the smallest valid graph. However it is not really practical to evaluate *all* the possible subcircuits, so we develop some heuristics to reduce the search space. Let us start by defining a few bounds on the considered subcircuits, we will leave the other ones apart in our algorithm:

- The subcircuits have at most B inputs ($\forall i, l^{(i)} < B$). The purpose of this bound is to limit the running time of Algorithm 1. In practice we take $B = 10$.
- The subcircuits are evaluable with one single bootstrapping with a maximum value p_{max} . This value ensures a bootstrapping with a reasonable timing. If the search algorithm fails for p_{max} , the subcircuit is dropped without trying to extend p . In practice, we take $p_{max} = 31$.

It could occur that several subcircuits share a common input, possibly encoded with different p -encodings for their homomorphic evaluation with our method. In practice this is not a problem since thanks to Lemma 12 one can change encoding under the same p “for free”. This justifies the fact that we can search for solution for each f_i independently and without caring about how they are connected.

In order to decompose our Boolean circuit into a graph satisfying the above property for a modulus p , we would want to exhaustively search all the subcircuits of \mathcal{F} compliant with the bounds we introduced earlier. However, all subcircuits are not as worth evaluating as the others. In particular a wire incoming a copy gate is particularly worth evaluating because it costs one bootstrapping but produce several inputs for the next subcircuits. Thus, our heuristic only considers the evaluation of subcircuits whose outputs incoming a copy gate, or the output z of the main graph \mathcal{F} .

Our heuristic works as follows:

1. For each of these outputs z_i , we exhaustively construct a set $\widehat{\mathcal{F}}_{z_i}$ that gathers all the evaluable subcircuits whose output is z_i and whose inputs are from previous copy gates. (Calling a subcircuit surrounded by copy gates an “atomic subcircuit”, the set $\widehat{\mathcal{F}}_{z_i}$ is composed of all the merges of atomic subcircuits that respect the global circuit wiring and have z_i as output). We then filter out the subcircuits of $\widehat{\mathcal{F}}_{z_i}$ that do not comply with the bounds introduced at the beginning of the section or that are not evaluable with the input modulus p (we use Algorithm 1 to decide this). If after filtering, the set $\widehat{\mathcal{F}}_{z_i}$ is empty, then our heuristic fails for the input value of p . At the end of this first phase, we have a set $\widehat{\mathcal{F}}$ containing all the candidate subcircuits that can be a node of a valid graph.

2. Now we want to construct the smallest valid graph evaluating \mathcal{F} using subcircuits from $\widehat{\mathcal{F}}$. While finding the smallest graph is hard, constructing any valid graph is easy. As a consequence, our strategy to find a small graph is to randomly create a lot of valid graphs and to take the smallest one. The procedure to create a valid graph is the following: we start from the output z and we randomly draw a subcircuit \mathcal{F}_z from $\widehat{\mathcal{F}}$ whose output is z . The inputs of \mathcal{F}_z can be sorted into two categories: the ones that are inputs of the whole circuit \mathcal{F} and the others, that are necessarily internal wires of \mathcal{F} . For each one of these latter wires w , we repeat the procedure, i.e. we draw a subcircuit \mathcal{F}_w from $\widehat{\mathcal{F}}$ whose output is w , and so on. When we have reached all the input wires of \mathcal{F} , we get a valid graph \mathcal{G} . This second step is run a large amount of times (the number of trials is a parameter of the method), and the smallest graph, i.e. the one with the fewest subcircuits, is returned.

5.3 Parallelization of the Execution of the Graph

Once we have our graph \mathcal{G} , we can identify its $n_{\mathcal{L}}$ layers. Formally, they are defined as:

Definition 14. A layer \mathcal{L} of a graph \mathcal{G} is a set of subcircuit $\{\mathcal{F}_\alpha, \dots, \mathcal{F}_\omega\}$ of \mathcal{G} that verifies: $\forall \mathcal{F}_i, \mathcal{F}_j \in \mathcal{L}, \mathcal{F}_i$ is not an ancestor node of \mathcal{F}_j .

By construction, all the subcircuits belonging to the same layer can be evaluated in parallel. This reduces the number of bootstrapping steps from k (the number of subcircuits in the graph \mathcal{G}) to $n_{\mathcal{L}}$ (the number of layers). Our graph-finding heuristic can be tweaked to select the graph with minimum number of layers instead of minimum number of subcircuits to optimize parallelization.

6 Adaptation of TFHE and the tfhe-rs Library

From a high level point of view, our technique can be seen as adding an additional layer of abstraction on top of TFHE. However things are not that simple: we had to tweak the bootstrapping procedure to make it work with odd values of p . Moreover, we implemented our framework using the `tfhe-rs` library [27] written in Rust. The following section covers the work of adaptation of the cryptographic scheme. The adaptation of the library is treated in Section 6.3.

6.1 Dealing with the Negacyclicity Problem for an Odd p

In the following, we explain the negacyclicity problem and how we intend to solve it. To do so, we need to dig into the details of the `BlindRotation` step of the bootstrapping, that we have introduced in Section 2.5.

Let $v(X)$ be a polynomial of the ring $\mathbb{Z}_{q,N}[X]/(X^N + 1)$, denoted by $v(X) = \sum_{k=0}^{N-1} v_k X^k$. Observe that a multiplication by X in this ring “rotates” the coefficients of the polynomial:

$$X \cdot v(X) = -v_{N-1} + v_0 \cdot X \cdots + v_{N-2} X^{N-1} .$$

In TFHE, the polynomial multiplication in the blind rotation is actually done by $X^{-\tilde{\mu}}$, with $\tilde{\mu} = \left\lfloor \frac{\mu \cdot 2N}{q} \right\rfloor$. This value lives in $\{0, \dots, 2N - 1\}$. This leads to two problems:

- A coefficient v_j can be brought in first place by two different rotations: the one induced by the polynomial multiplication by $X^{[-j]_{2N}}$ and the one by $X^{[-j+N]_{2N}}$.
- Each time a coefficient goes last to first, it gets negated (because $X^N = -1$ in the ring). So actually, the multiplication by $X^{[-j]_{2N}}$ yields correctly v_j , but the one by $X^{[-j+N]_{2N}}$ yields $-v_j$.

However, these problems are not equally serious for even and odd values of p . Recall that $\mu = m + e \in \mathbb{Z}_q$, with e sampled from a small centered Gaussian. The use of a small error makes that μ does not take all the values of \mathbb{Z}_q with the same probability: in particular, the densest parts in terms of probability over \mathbb{Z}_q



Fig. 5: Distribution of the values of μ across \mathbb{Z}_q for $p = 6$ and $p = 5$: the colored parts show the dense spots where the value has a high probability to lie in. The width of these sectors depends on σ (the standard deviation of the error distribution χ of TFHE). Note that this repartition looks the same for $\tilde{\mu}$ in \mathbb{Z}_{2N} .

are the one close to the “unscrambled” values of m , namely $\left\{ \frac{kq}{p} \mid k \in \mathbb{Z}_p \right\}$ (in the following, we ignore the rounding operation necessary if q is not divisible by p). We illustrate this distribution on Figure 5. We call these sections of the torus the *dense spots*.

When we transpose these dense spots into \mathbb{Z}_{2N} , they become the sectors close to $\left\{ \frac{k \cdot 2N}{p} \mid k \in \mathbb{Z}_p \right\}$ (we keep ignoring the rounding). Let us note that the noises in \mathbb{Z}_q and \mathbb{Z}_{2N} are fundamentally different: the former is the one added at encryption that may have grew during the homomorphic computations, and the latter is called “drift” and is caused by the accumulation of the rounding errors on each coefficient of the ciphertext during the modulus switching (but this difference in nature does not impact our purpose). Let $k \in \mathbb{Z}_p$, the multiplication $X^{-\frac{k \cdot 2N}{p}} \cdot v(X)$ yields the same degree-zero coefficient as the multiplication $X^{\left[-\frac{k \cdot 2N}{p} + N\right]_{2N}} \cdot v(X)$, up to the minus sign. For the sake of clarity, we write the exponent of the latter in a slightly different manner:

$$\left[\frac{-k \cdot 2N}{p} + N \right]_{2N} = \left[\frac{(-k + \frac{p}{2}) \cdot 2N}{p} \right]_{2N}$$

This is where the parity of p plays a part: if p is even, then $\left[\frac{(-k + \frac{p}{2}) \cdot 2N}{p} \right]_{2N}$ is a dense spot as well. So, the rotations by these two values will happen with high probability and they will both yield the same coefficient $v_{\frac{k \cdot 2N}{p}}$ (up to the minus sign for one of them). Thus, when evaluating a function f with a PBS, the calls $f(k)$ and $f(k + \frac{p}{2})$ will produce the same output (one again, up to the minus sign), which is a collision constraining the definition of f . On the other hand, let us consider an odd value to p . Then, $\left[\frac{(-k + \frac{p}{2}) \cdot 2N}{p} \right]_{2N}$ is no longer a dense spot, as it lies exactly halfway between the two dense spots $\left[\frac{(-k + \frac{p-1}{2}) \cdot 2N}{p} \right]_{2N}$ and $\left[\frac{(-k + \frac{p+1}{2}) \cdot 2N}{p} \right]_{2N}$. As a consequence, collision never occurs. Figure 6 illustrates this phenomenon.

As a consequence, we select only odd values for p in our framework. The downside is that the acceptable bound on the error level in the ciphertexts is tighter. We mentioned earlier in Section 2.3 that in the original TFHE scheme the bound on error for correct decryption was $e < \frac{q}{2p}$. In the odd setup, this bound is divided by 2, so we get $e < \frac{q}{4p}$. We will see in Section 6.4 how this change impacts the parametrization of the scheme.

Exception for $p = 2$: We just said that only odd values can be selected for p in our framework, however p -encodings with even values of p exist as well: nonetheless they need to achieve the relaxed negacyclicity property introduced in Definition 5. This restriction makes them basically useless, as using only odd p -encodings is sufficient to evaluate all possible Boolean functions without having to bother with the negacyclicity property. However, the case $p = 2$ is an exception: the valid 2-encodings are automatically negacyclic



(a) With p even, the dense spots of each half of the torus are aligned. (b) With p odd, the dense spots face empty spots, close to the bounds of the p -sectors.

Fig. 6

and allow the PBS-free evaluation of series of XOR operations. So it might be efficient to switch between 2-encodings for XOR operations and p -encodings (with odd p) for non-linear Boolean functions. We make use of this strategy in our implementation of the Keccak permutation in Section 7.3.

6.2 Construction of the Accumulator for an Odd p

The accumulator is the polynomial $v(X)$ used in the `BlindRotate` step of the bootstrapping. In the Section 6.1, we showed how the values are spread over the torus after bootstrapping. To implement this, we need to explicitly characterize this polynomial. In the following presentation, we neglect roundings to keep notations light (as if p would divide N), or, equivalently, the division operator is assumed to include rounding.


Definition 15. *If p is an odd modulus, and $f : \mathbb{Z}_p \mapsto \mathbb{Z}_{p'}$ a function, then the accumulator $v(X) \in \mathbb{Z}_{N,q}[X]/(X^N + 1)$ has the form:*

$$v(X) = X^{-\frac{N}{2p}} \cdot \sum_{j=0}^{N/p-1} X^j \cdot \left(\sum_{i=0}^{\frac{p-1}{2}} f(i) X^{i \frac{2N}{p}} + \sum_{i=0}^{\frac{p-1}{2}-1} -f\left(i + \frac{p+1}{2}\right) X^{i \frac{2N}{p} + \frac{N}{p}} \right)$$

In the following, we explain the structure of this accumulator. The polynomial has degree N and is made of p distinct windows of width $\frac{N}{p}$. Each of these windows has constant coefficient value $f(k)$, for $k \in \{0, \dots, p-1\}$. For $0 \leq \alpha \leq \frac{p-1}{2}$, the range of degrees whose coefficients are $f(\alpha)$ is $\left[\alpha \frac{2N}{p} - \frac{N}{2p}; \alpha \frac{2N}{p} + \frac{N}{2p}\right]$. Now, for $\frac{p+1}{2} \leq \beta \leq p-1$, we can write $\beta = \alpha + \frac{p+1}{2}$, with $0 \leq \alpha < \frac{p-1}{2}$. This time, the range of spanned degrees is $\left[\alpha \frac{2N}{p} + \frac{N}{2p}; (\alpha+1) \frac{2N}{p} - \frac{N}{2p}\right]$. Thus, the values $k \in \{0, \dots, p-1\}$ spans the entire space $[0; N)$ without overlap. The values over $\frac{p+1}{2}$ gets negated by the negacyclicity, so the underlying coefficient is also negated to compensate this effect. We illustrate this construction on Figure 7.

6.3 Concrete Implementations of p -Encodings and Homomorphic Functions in `tfhe-rs`

To implement our framework, we relied on the `tfhe-rs` library [27]. Here is a list of the major changes we applied to the code:



$$\begin{array}{l}
 v(X) = \begin{array}{|c|c|c|c|c|c|} \hline f(0) & -f(3) & f(1) & -f(4) & f(2) & -f(0) \\ \hline \end{array} \\
 X^N \cdot v(X) = \begin{array}{|c|c|c|c|c|c|} \hline -f(0) & f(3) & -f(1) & f(4) & -f(2) & f(5) \\ \hline \end{array} \\
 \begin{array}{cccccc}
 0 & \frac{N}{2p} & \frac{N}{p} & \frac{2N}{p} & \frac{3N}{p} & \frac{4N}{p} & N \\
 \end{array}
 \end{array}$$

Fig. 7: Illustration of the construction of the accumulator. On top is the ring \mathbb{Z}_{2N} splitted in windows. Below is a representation of the polynomial v , with its version once rotated by a multiplication by X^N .

Addition of the notion of p -encoding: An encoding \mathcal{E} is simply implemented with a structure **Encoding** storing two **HashSets** and the modulus p . The **HashSets** represent both sets $\mathcal{E}(0)$ and $\mathcal{E}(1)$. When creating an **Encoding**, the software checks whether the two underlying sets are disjoint or not. Moreover, the operation of encryption and decryption are modified as well. The signatures change from:

`encrypt(Boolean, ClientKey) -> Ciphertext`

to:

`encrypt(Boolean, ClientKey, Encoding) -> Ciphertext`

(same for `decrypt`). The functions also perform the mapping $\mathbb{B} \mapsto \mathbb{Z}_p$ before encryption and the other way around after decryption.

*Definition of the new structure **HomFunc**:* According to the evaluation strategy we introduced in Section 3.2, we wrote a new structure **HomFunc**, associated to a Boolean function $f : \mathbb{B}^\ell \mapsto \mathbb{B}$, carrying:

- A list of the **Encoding** objects for the inputs: $\mathcal{E}_{in} = (\mathcal{E}_1, \dots, \mathcal{E}_\ell)$, with the input modulus p_{in} they encoded on.
- The output **Encoding** object \mathcal{E}_{out} , with the output modulus p_{out} it is encoded on.
- The clear function f .

When such a structure is constructed, it self-checks whether $f(\mathcal{E}_{in})$ is valid. Then, when provided ℓ **Ciphertexts** objects encoded under their respective p -encoding, it executes the homomorphic sum and the PBS and outputs the results encoded under \mathcal{E}_{out} . Some utility functions performing encoding-switching are also available, allowing the chaining of several **HomFunc**.

Implementation of the accumulator: The procedure of bootstrapping of **tfhe-rs** is slightly modified to support the new version of the accumulator we introduced in Section 6.2.

Parsing of graphs: We implemented a Python script that produces graphs to represent more complex functions that requires several PBS, as described in Section 5. These graphs are stored with a comprehensive file format and our Rust implementation has a module of parsing allowing to load these graphs and automatically generate the corresponding graph of **HomFunc**.

6.4 Crafting of Parameters

The instances of the TFHE scheme are defined by a set of parameters. These parameters should simultaneously ensure the security of the scheme and the correctness of the homomorphic computations. The security of the LWE instances can be estimated with the `lattice-estimator` [1] that gathers all the known attacks on lattice-based schemes and gives an estimation of their level of security. The correctness of the scheme depends on the statistical variance of the noise during the computation, which should not exceed some critical bound.

Here is a concrete overview of the different parameters at play in an instance of TFHE:

- n : the dimension of the LWE samples. Namely, the TLWE ciphertexts are vectors of length $n + 1$.
- q : the modulus of the ring the encrypted values live on. In `tfhe-rs` those values are stored on `u32` values, making $q = 2^{32}$. We treat this as an immutable value.
- σ : the standard deviation of the Gaussian distribution of error in LWE samples.
- k : the dimension of the GLWE samples. Namely, the TRLWE samples are vectors of polynomial of length $k + 1$. In practice, we consider only cases where $k = 1$ (actually making those RLWE samples).
- σ' : the standard deviation of the Gaussian distribution of error in GLWE samples.
- A few more parameters dimensioning some inner algorithms of the bootstrapping. A detailed description and an analysis of their impact on performances and noise level can be found in [5]. In this work, they are denoted as *micro-parameters*.

Following the work of [5], the tool [26] generating appropriate sets of parameters has been developed. It takes as inputs the number of bits of precision (in our case, the number of bits of the modulus p) and the sum of the squares of the coefficients whose ciphertexts are multiplied with. In our case, it is simply the 2-norm of the vector \mathbf{d} computed by the Algorithm 1, which we denote by $\|\mathbf{d}\|$. Our protocol to select parameters suitable with our framework is to run the tool for the precision $n_b = \lceil \log_2(p) \rceil + 1$ bits (where the extra bit is added because the bound on tolerable error is divided by 2 compared to the original TFHE).

7 Application to Cryptographic Primitives

In this section, we apply our approach on some cryptographic primitives. For each primitive, we first explain the regularization of the Boolean function, then the dimensioning of the cryptographic parameters, and finally report the concrete performances of our implementation. We detailed all the timings of our experimentations along with the sets of parameters we used in Section 7.6. The code of these experimentations can be found at <https://github.com/CryptoExperts/bpr-boolean-fhe>.

For performance measurement, we implemented our framework in the library `tfhe-rs` [27] adapted as discussed in Section 6 we and generated the sets of parameters thank to `concrete-optimizer` [26]. All experiments have been carried out on a laptop with a 12th Gen Intel(R) Core(TM) i5-1245U CPU with 10 cores and a frequency of 4.4 GHz, and 16 GB of RAM.

7.1 SIMON Block Cipher

SIMON is a hardware-oriented block cipher developed in [3], which relies only on the following operations: AND, rotation, XOR. It is a classical Feistel network for which the Feistel function consists in applying basic operations on the branch, xoring the subkey and then xoring the result with the other branch as depicted in the Figure 8. We use one ciphertext per bit so the rotation operation is essentially free. Note that the key is considered as a plaintext, which does not change anything in the framework. In our implementation, we considered a (128-128) instance of SIMON (i.e. the whole state and the key are of size 128).

The Boolean function to evaluate can be defined as

$$f(b_0, b_1, b_2, b_3, b_4) = b_0 \cdot b_1 \oplus b_2 \oplus b_3 \oplus b_4 .$$




Fig. 8: One Feistel round of SIMON. S^i denotes the left circular shift by i bits. This figure is extracted from the original paper [3].

Using Algorithm 1, we found the smallest possible p ($p = 9$) and the following 9-encodings to evaluate each bit of the Feistel function with one single bootstrapping (i.e. totalling 64 PBS per round).

$$\mathcal{E}_0 = \mathcal{E}_1 = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{1\} \end{cases} \quad \text{and} \quad \mathcal{E}_2 = \mathcal{E}_3 = \mathcal{E}_4 = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{2\} \end{cases} \quad \text{with } p = 9.$$

The sum of these p -encodings yields the output encoding:

$$\mathcal{E}_{out} = \begin{cases} 0 \mapsto \{0, 1, 4, 5, 8\} \\ 1 \mapsto \{2, 3, 6, 7\} \end{cases} \quad \text{with } p = 9$$

which is valid for f . After the PBS, all the bits of the state are encrypted under the encoding \mathcal{E}_0 .

To perform a Feistel round on a state of size k , the function f is applied in parallel $k/2$ times. Note that one bit may be used in several evaluation as b_0 , b_1 and b_2 . So we sometimes have to switch from \mathcal{E}_0 to \mathcal{E}_1 by a simple external multiplication by 2, which is negligible in terms of performances.

Using the `concrete-optimizer` [26], we crafted a set of parameters suitable for this modulus. On our machine, one PBS with such parameters takes about 11 ms. The theoretical timings achieved on one full block without any parallelization is 47 seconds (68 rounds \times 64 bits \times 11 ms) which we confirmed experimentally.

Nonetheless, this setting is intrinsically parallelizable: the 64 bootstrappings of each round can be performed in parallel. We implemented parallelization using the module `Rayon` of Rust, which made the total timings drop to 13 seconds on our machine.

Compared to [4] that implemented the same block cipher on an equivalent hardware with parallelism, our implementation is about 10 times faster. Table 5 shows the comparison.

7.2 The Trivium Stream Cipher

Trivium [9] is a stream cipher that uses a circular state. At each round, the bits are rotated within the state, except for three of them that are refreshed using the Boolean function of Section 7.1. The outer stream is generated by xoring three bits of the state each round once a “warming-up” phase is achieved.

For each generated key bit, it requires performing this function three times and aggregating five XOR operations in the center. Our strategy is to evaluate the refreshing function three times per round with one

PBS for each of them, then get the result in \mathbb{Z}_2 and chain the five XOR operations to get the output. The Figure 9 illustrates the layout of the cipher.




Fig. 9: The trivium stream cipher. Figure extracted from [9]

In [2], the authors implement Trivium using the original `tfhe-rs` library, with using two bits of message and 2 bits of carry for a total of 4 significative bits out of the 32 of a ciphertext component. They call this mode the `shortint` mode. The use-case they target is transciphering.

To compare our implementation with the one of [2], timings are not a good metric as in their work they are provided on a massive AWS instance with a significant amount of parallelism. A better metric is to count the number of PBS and compare the parameter sets.

We reproduced the PBS operation with their parameter set on our machine and then simply estimated the timings of one round of Trivium with their approach with no parallelism. The results are summed up in Table 1. Note that in our implementation we do not refresh the output bits with a PBS after the chain of XOR, because in the use-case of transciphering one more XOR has to be performed with the message. We take advantage of this and move the last PBS in the transciphering phase.

Table 1: Comparison of timings of one round of Trivium between our work and [2]

Instance	Timing PBS	Number of PBS per round	Estimated timings
[2]	6.6	7	46.2 ms
Our work	11	3	33 ms

7.3 Keccak Permutation

Keccak is a hash function standardized by NIST under the name *SHA-3* [23]. It is a sponge function, whose transformation is called the *Keccak permutation*. It consists of five sub-functions: θ , ρ , π , χ , and ι .

Let us recall that our approach encrypts each bit in one TFHE ciphertext. Let us explain the strategies of homomorphization of these sub-functions:

- ρ and π simply reorder the bits within the state, so they are not impacted by the homomorphization.
- θ is just a serie of XOR operations, so it can be performed with a serie of homomorphic additions and without any PBS provided that the input ciphertexts are defined over \mathbb{Z}_p with $p = 2$.
- χ is the only non-linear function of the permutation, and has to be performed with a PBS. It is the transformation that applies the function defined by

$$f_\chi(a, b, c) = a \oplus c \oplus b \& c$$

to get each bit of the output state.

- Finally, ι performs a simple xor with a constant, so it can be handled in a similar manner that θ . The difference is that the constant is in clear this time.

The p -encodings we use are:

- $\mathcal{E}_\& = \begin{cases} 0 \mapsto \{1\} \\ 1 \mapsto \{2\} \end{cases}$ with $p_\& = 3$ to evaluate the $\&$ operator in the alternative formula of χ .
- $\mathcal{E}_\oplus = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{1\} \end{cases}$ with $p_\oplus = 2$ for the other operations of \oplus .

Our strategy of homomorphic evaluation of the Keccak permutation is as follows:

1. Encrypt the input state under the encoding \mathcal{E}_\oplus .
2. Evaluate the subfuctions θ , ρ , and π .
3. Change the encoding from \mathcal{E}_\oplus to $\mathcal{E}_\&$ with one PBS per bit of the state.
4. Evaluate the AND operator of the subfunction χ and change the encoding from $\mathcal{E}_\&$ to \mathcal{E}_\oplus with one PBS per bit of the state.
5. Evaluate the remaining \oplus operators of χ and the ι subfunctions, then restart from Step 1. for the next loop iteration.

While converting ciphertexts from $\mathcal{E}_\&$ to \mathcal{E}_\oplus is trivial with a PBS, the other way around is more tricky because p_\oplus is even. Because of the negacyclicity problem, to be able to perform the encoding switching $\mathcal{E}_\& \rightarrow \mathcal{E}_\oplus$, one needs $\mathcal{E}_\&(0) = [-\mathcal{E}_\&(1)]_{p_\&}$. With $p_\& = 3$, the only candidate is the encoding $\mathcal{E}_\&$ defined above.

As a result, each round takes two programmable bootstrappings per bit. An implementation with our tweaked version of `tfhe-rs` takes 22 seconds (without any parallelism) on our hardware to perform one Keccak round on a state of 1600 bits in spite of the two PBS required per round and per bit. Those timings are possible because of the small values of p allowing the use of small sets of parameters, which speeds up the computation. A full run of Keccak counting 24 rounds, we can then estimate the timings to 8.8 minutes. For the sake of simplicity, we use the same set of parameters for both types of PBS.

7.4 Ascon

Ascon [17] is a lightweight block cipher algorithm that was designed to provide efficient and secure encryption and authentication for a wide range of applications, particularly in resource-constrained environments such as embedded systems and IoT devices. The name ‘‘Ascon’’ stands for ‘‘Authenticated encryption for Small Constrained Devices’’. We implemented its s-box, whose circuit is represented on Figure 10.

This layout is a bit different from the others: we focus on the 5-bit to 5-bit s-box and we denote f_0, \dots, f_4 the five functions of $\mathbb{B}^5 \mapsto \mathbb{B}$ that generate the 5 output bits x_0, \dots, x_4 .

These functions, once analyzed by the algorithm, can be computed in one single bootstrapping each, but for different values of p (respectively $p = 17, 7, 7, 15, 11$ that are the smallest possible values). To make things




Fig. 10: The 5-bits look-up table of ASCON. Figure extracted from [17]

simpler, we generated only encodings with $p = 17$, making the implementation more straightforward with a little loss of performances. For each subfunction f_i , 5 canonical 17-encodings $(\mathcal{E}_{i,0}, \dots, \mathcal{E}_{i,4})$ of form

$$\mathcal{E}_{i,j} = \begin{cases} 0 \mapsto \{0\} \\ 1 \mapsto \{d_{i,j}\} \end{cases}$$

are computed. The results are displayed in the Table 2. Note the zero values in some cases, they show that the variable is not used in the subfunction.

subfunction	$d_{i,0}$	$d_{i,1}$	$d_{i,2}$	$d_{i,3}$	$d_{i,4}$
f_0	1	2	3	7	14
f_1	1	2	2	2	4
f_2	1	2	2	4	0
f_3	1	1	5	5	3
f_4	1	2	0	4	3

Table 2: Parameters $d_{i,j}$ for Ascon, with $p = 17$ for every subfunction.

Using this solution, the s-box is evaluated in 92 ms. Note that the 5 different PBS described in Table 2 have different norms of vector \mathbf{d} so they may have a different set of parameters for each. For the sake of simplicity, we use the more restrictive one for the 5. Estimating the timings of a full run of Ascon is not trivial because it depends a lot of the parameters. To give a rough idea, in hashing mode, 64 s-boxes are required per round, with 12 rounds recommended. The outputs of the s-boxes are in \mathbb{Z}_2 to allow the evaluation of the linear layer of Ascon. At the end of this linear layer, the encoding of each of the 320 bits of the state must be switched back to \mathbb{Z}_{17} with a PBS. To do so, we use the same set of parameters as for the encoding switching in Step 3 of the Keccak evaluation in Section 7.3.

This gives a rough estimation of 92.16 seconds for one Ascon hash.

7.5 AES

AES [16], or Advanced Encryption Standard, stands as one of the most widely used and trusted encryption algorithms in the world of computer security. Its standardization occurred in 2001 when it was adopted by NIST to replace the obsolete DES (Data Encryption Standard). Implementing this primitive in FHE is known as particularly tricky and only few attempts have been made [19], [15], [25].

A round of AES can be decomposed into 4 steps:

1. **SubBytes**: a non-linear substitution step where each byte is replaced by another according to a lookup table. This step concentrates all the challenge for homomorphization, the other one being trivial with our framework.
2. **ShiftRows**: a transposition step where the last three rows of the state are shifted cyclically a certain number of times. As our framework encrypts each bit in a distinct ciphertext, this step is for free.
3. **MixColumns**: a linear mixing operation which operates on the columns of the state, combining the four bytes in each column. This step can be implemented using only XOR operations and bit-shiftings. The former are trivial with our framework using $p = 2$ and the latter are for free as the ones in the previous step.
4. **AddRoundKey**: each byte of the state is combined with a byte of the key from the key schedule using a XOR. Still using $p = 2$, this can be carried out easily.

Thus, we focus here on the implementation of the s-box of the **SubBytes** step, also known as the *Rijndael s-box*. This s-box takes 8 bits in input and yields 8 bits of output. It is defined by two substeps: an inversion in $GF(2^8)$ followed by an affine transformation. While the latter is trivial to compute with TFHE, the former is much trickier and thus we did not take advantage of this representation.

Using our framework, the obvious-looking solution is to split the full s-box $\mathbb{B}^8 \mapsto \mathbb{B}^8$ into 8 subfunctions $f_0, \dots, f_7 : \mathbb{B}^8 \mapsto \mathbb{B}$. We could then give them to the search algorithm of Section 4. If this would work, we could evaluate the Rijndael s-box in 8 PBS. Unfortunately, the algorithm does not converge for values of p “reasonable”.

We thus need to leverage an alternative representation of the s-box. A well known efficient Boolean representation of the AES s-box is given in [6]. In this work, authors applied logic minimization techniques to produce an optimized Boolean circuit (in terms of number of gates) of the s-box splitted in 3 phases:

1. A purely linear layer mapping the 8 input bits onto 22 bits.
2. A middle non-linear layer, represented by a circuit with exclusively AND and XOR logic gates, mapping the previous 22 bits onto 18 bits.
3. A final purely linear layer mapping the 18 bits on the 8 output bits of the s-box.

In the following, we focus on the implementation of the second step, the two others being trivial with TFHE. Of course, we cannot split the layer into 18 subfunctions $f : \mathbb{B}^{22} \rightarrow \mathbb{B}$ and run our algorithm of search on every single one of them, as the runtime would explode due to their high arity. Instead, we analyze the Boolean circuit and apply the approach introduced in Section 5.

Homomorphization of the S-box. We start from the circuit representation given in the work of [6]. This set of instructions is compiled into a circuit \mathcal{A} , compliant with the definitions introduced in Section 5.1.

Each of the 18 outputs (z_0, \dots, z_{17}) are isolated from each other and the circuits ($\mathcal{A}_0, \dots, \mathcal{A}_{17}$) generating them are separated. Of course, some intermediary values are used in several circuits, but for now we ignore this and we considerate the 18 problems as independent from each other.

Then, for each circuit \mathcal{A}_i , we create the sets $\hat{\mathcal{G}}_i$ containing respectively all the valid subcircuits of \mathcal{A}_i , as introduced in Section 5.1. As the circuits are intricated with each others, these sets share a large numbers of subcircuits in common, so they are computed all together at the same time. A small graph in each set is computed using the strategy developed in Section 5.2.

Finally we merge all the graphs and run everything for a total of 36 PBS to evaluate the full circuit \mathcal{A} , with a global $p = 11$. This allows a relatively quick bootstrapping.

Recall that the **SubBytes** step is made of 16 s-boxes. So, we can extrapolate that one execution of the **SubBytes** step would take $16 \times 36 = 576$ PBS. The outputs of this step would be encoded with $p = 2$, allowing the XOR operations of the following steps to be performed efficiently. Namely, **MixColumns** can be reduced to a serie of XOR and bit shifts, and **AddRoundKey** is basically one XOR per bit of the state. We also need to take into account the encoding switching to come back to $p = 11$ before each **SubBytes**. It costs one PBS per bit, so 128 PBS. Finally, this gives a total of 704 PBS per round. For AES-128, which takes 10 rounds, we estimate a full run to 7040 PBS.

Performances. In terms of performances, with the set of parameters we selected, a PBS takes 30 ms on our hardware. We implemented the non-linear layer of the **SubBytes** steps (that concentrates the majority of the PBS) and an evaluation takes 1.1 second (which is consistent with the 36 PBS) without any parallelization. To get the estimated runtime for a full **AES-128**, we will need $7040 \times 0.030 = 211$ seconds without any parallelization. We note that the 16 evaluations of s-boxes in **SubBytes** can be parallelized, as well as each of the 128 encoding switchings before **SubBytes**. Moreover, within each s-box, we can locally apply our strategy of parallelization introduced in Section 5.3. On our machine, this intra s-box parallelization makes the time of evaluation of the 8-bits s-box drop from 1.1 s to 450 ms.

We compare favorably to previous works of [19] and [15], who report timings of respectively 18 minutes and 5 minutes for a full AES. The more recent work of [25], also proposes an implementation of **AES-128** using a completely different technique called the *tree-bootstrapping*. On a similar experimental setup, they claim an execution in 270s that we slightly beat with our 211s run.

7.6 Summary of Applications

We summarize hereafter the parameters and performances of our implementations of cryptographic primitives. Table 3 gives an overview of the TFHE parameters used for each value of p in these examples. They all meet the required level of security of 2^{128} . It also shows the associated p and the norm of \mathbf{d} , denoted by $\mathcal{N}_{\mathbf{d}}$ (that corresponds to $\mathcal{N}_{\mathbf{d}} = \lceil \log_2(\|\mathbf{d}\|) \rceil$) that are the input of the parameter selection algorithm. Table 4 shows the complexity of the cryptographic primitives expressed in PBS. Finally, Table 5 shows the concrete performance achieved by our implementations on our machine, as well as the comparison with other works. For more information about an implementation or a comparison, the reader is referred to the related section.

Identification		Constraints		TFHE parameters									Timings
Ref.	Sections	p	$\mathcal{N}_{\mathbf{d}}$	n	k	N	σ_{LWE}	σ_{GLWE}	B_g	ℓ_g	B_{KS}	t_{KS}	PBS
$\text{PBS}_{(9,2)}$	7.1, 7.2	9	2	774	1	2048	$2.15e^{-6}$	$2.17e^{-19}$	23	1	4	3	11 ms
$\text{PBS}_{(3,2)}$	7.3	3	2	668	6	256	$2.04e^{-5}$	$3.45e^{-12}$	18	1	4	3	7 ms
$\text{PBS}_{(2,1)}$	7.3, 7.4	3	2	668	6	256	$2.04e^{-5}$	$3.45e^{-12}$	18	1	4	3	7 ms
$\text{PBS}_{(17,5)}$	7.4	17	5	768	1	2048	$7.07e^{-6}$	$2.94e^{-16}$	15	2	5	3	17 ms
$\text{PBS}_{(11,6)}$	7.5	11	6	807	1	4096	$2.15e^{-6}$	$2.16e^{-19}$	22	1	5	3	30 ms

Table 3: Sets of TFHE parameters for each PBS used in our implementations, with the constraints used to generate the sets, and the performances. Each setting is referenced as $\text{PBS}_{(p,\ell)}$ with $\ell = \lceil \log_2(\|\mathbf{d}\|) \rceil$.

Section	Primitive	Complexity in PBS
7.1	One round of SIMON-128	64 $\text{PBS}_{(9,2)}$
	One full run of SIMON-128	4352 $\text{PBS}_{(9,2)}$
7.2	One round of Trivium	3 $\text{PBS}_{(9,2)}$
	One warm-up phase of Trivium (*)	3456 $\text{PBS}_{(9,2)}$
7.3	One round of Keccak	1600 $\text{PBS}_{(3,2)}$ + 1600 $\text{PBS}_{(2,1)}$
	A full Keccak permutation (*)	38400 $\text{PBS}_{(3,2)}$ + 38400 $\text{PBS}_{(2,1)}$
7.4	One evaluation of Ascon’s s-box	5 $\text{PBS}_{(17,5)}$
	One full Ascon hashing run (*)	3840 $\text{PBS}_{(17,5)}$ + 3840 $\text{PBS}_{(2,1)}$
7.5	One evaluation of the AES s-box	36 $\text{PBS}_{(11,6)}$
	A full run of AES-128 (*)	5760 $\text{PBS}_{(11,6)}$ + 1280 $\text{PBS}_{(2,1)}$

Table 4: Complexity of the different primitives we implemented with respect to the PBS of Table 3. The primitives indicated by a (*) are estimations while the others have been fully implemented.

Primitive	Section or Other work	Performances
One full run of SIMON	[4]	128 s
	Our work (Section 7.1)	13 s
One warm-up phase of Trivium (*)	[2] (estimation on our machine)	53 s
	Our work (Section 7.2)	38 s
One Full Keccak permutation (*)	Our work (Section 7.3)	8.8 min
One Ascon hashing (*)	Our work (Section 7.4)	92 s
One full evaluation of AES-128 (*)	[19]	18 min
	[15]	5 min
	[25]	270 s
	Our work (Section 7.5)	211 s

Table 5: Timings of evaluation of full primitives, and comparison with previous works when they exist. Like on Table 4, a star (*) is added in the cells if the timings is not obtained from a full implementation but estimated from an implemented building block.

8 Conclusion

In this paper, we have proposed a new strategy to evaluate Boolean functions homomorphically using TFHE. Our technique relies on constructing an intermediate homomorphic layer between the Boolean space \mathbb{B} of the plaintexts and the torus \mathbb{T}_q on which ciphertexts live. We introduced a formal model for our technique and detailed algorithms to efficiently construct such layers. We further extended our strategy to the case of arbitrary Boolean circuits by developing some heuristic to decompose a circuit into Boolean functions efficiently evaluable with our framework. We applied our framework to various cryptographic primitives, in particular to the challenging AES cipher. All the reported implementations outperform the state of the art.

Future work may focus on a thinner analysis of the noise behaviour to develop a better protocol to select parameters. The search algorithm can probably be enhanced to scale better with the arity of the input function. Finally, more work on the efficient decomposition of Boolean circuits would be welcome: especially if we want to evaluate deeper circuits.

References

1. Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, October 2015.
2. Thibault Balenbois, Jean-Baptiste Orfila, and Nigel P. Smart. Trivial transciphering with trivium and tfhe. Cryptology ePrint Archive, Paper 2023/980, 2023. <https://eprint.iacr.org/2023/980>.
3. Ray Beaulieu, Douglas Shors, Jason Smith, Stefan Treatman-Clark, Bryan Weeks, and Louis Wingers. The simon and speck families of lightweight block ciphers. Cryptology ePrint Archive, Paper 2013/404, 2013. <https://eprint.iacr.org/2013/404>.
4. Adda-Akram Bendoukha, Oana Stan, Renaud Sirdey, Nicolas Quero, and Luciano Freitas. Practical homomorphic evaluation of block-cipher-based hash functions with applications. Cryptology ePrint Archive, Paper 2023/480, 2023. <https://eprint.iacr.org/2023/480>.
5. Loris Bergerat, Anas Boudi, Quentin Bourgerie, Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Parameter optimization and larger precision for (t)he. Cryptology ePrint Archive, Paper 2022/704, 2022. <https://eprint.iacr.org/2022/704>.
6. Joan Boyar and Rene Peralta. New logic minimization techniques with applications to cryptology. Cryptology ePrint Archive, Paper 2009/191, 2009. <https://eprint.iacr.org/2009/191>.
7. Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical gapsvp. Cryptology ePrint Archive, Paper 2012/078, 2012. <https://eprint.iacr.org/2012/078>.
8. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)*, 6:1 – 36, 2014.

9. Christophe Cannière. Trivium: A stream cipher construction inspired by block cipher design principles. volume 4176, pages 171–186, 08 2006.
10. Hao Chen, Iliaria Chillotti, and Yongsoo Song. Improved bootstrapping for approximate homomorphic encryption. Cryptology ePrint Archive, Paper 2018/1043, 2018. <https://eprint.iacr.org/2018/1043>.
11. Jung Hee Cheon, KyooHyung Han, Andrey Kim, Miran Kim, and Yongsoo Song. Bootstrapping for approximate homomorphic encryption. Cryptology ePrint Archive, Paper 2018/153, 2018. <https://eprint.iacr.org/2018/153>.
12. Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. Homomorphic encryption for arithmetic of approximate numbers. Cryptology ePrint Archive, Paper 2016/421, 2016. <https://eprint.iacr.org/2016/421>.
13. Iliaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. Tfhe: Fast fully homomorphic encryption over the torus. Cryptology ePrint Archive, Paper 2018/421, 2018. <https://eprint.iacr.org/2018/421>.
14. Iliaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Improved programmable bootstrapping with larger precision and efficient arithmetic circuits for tfhe. Cryptology ePrint Archive, Paper 2021/729, 2021. <https://eprint.iacr.org/2021/729>.
15. Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Scale-invariant fully homomorphic encryption over the integers. Cryptology ePrint Archive, Paper 2014/032, 2014. <https://eprint.iacr.org/2014/032>.
16. Joan Daemen and Vincent Rijmen. The block cipher rijndael. In Jean-Jacques Quisquater and Bruce Schneier, editors, *Smart Card Research and Applications*, pages 277–284, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.
17. Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schläpfer. Ascon v1.2: Lightweight authenticated encryption and hashing. *J. Cryptol.*, 34(3):33, 2021.
18. Craig Gentry. Fully homomorphic encryption using ideal lattices. volume 9, pages 169–178, 05 2009.
19. Craig Gentry, Shai Halevi, and Nigel P. Smart. Homomorphic evaluation of the aes circuit. Cryptology ePrint Archive, Paper 2012/099, 2012. <https://eprint.iacr.org/2012/099>.
20. Marc Joye. Guide to fully homomorphic encryption over the [discretized] torus. Cryptology ePrint Archive, Paper 2021/1402, 2021. <https://eprint.iacr.org/2021/1402>.
21. Joon-Woo Lee, Eunsang Lee, Yongwoo Lee, Young-Sik Kim, and Jong-Seon No. High-precision bootstrapping of rns-ckks homomorphic encryption using optimal minimax polynomial approximation and inverse sine function. Cryptology ePrint Archive, Paper 2020/552, 2020. <https://eprint.iacr.org/2020/552>.
22. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. Cryptology ePrint Archive, Paper 2012/230, 2012. <https://eprint.iacr.org/2012/230>.
23. NIST. SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions, 2015. <https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.202.pdf>.
24. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), sep 2009.
25. Daphné Trama, Pierre-Emmanuel Clet, Aymen Boudguiga, and Renaud Sirdey. At last! a homomorphic aes evaluation in less than 30 seconds by means of tfhe. Cryptology ePrint Archive, Paper 2023/1020, 2023. <https://eprint.iacr.org/2023/1020>.
26. Zama. concrete-optimizer : Concrete Optimizer is a Rust library that find the best cryptographic parameters for a given TFHE homomorphic circuit., 2022. <https://github.com/zama-ai/concrete/tree/main/compiler/concrete-optimizer>.
27. Zama. TFHE-rs: A Pure Rust Implementation of the TFHE Scheme for Boolean and Integer Arithmetics Over Encrypted Data, 2022. <https://github.com/zama-ai/tfhe-rs>.