

TFHE Gets Real: an Efficient and Flexible Homomorphic Floating-Point Arithmetic

Loris Bergerat^{1,2}, Ilaria Chillotti^{*}, Damien Ligier^{*},
Jean-Baptiste Orfila¹, and Samuel Tap¹

¹Zama, Paris, France - <https://zama.ai/>

{loris.bergerat, jb.orfila, samuel.tap}@zama.ai

²Université Caen Normandie, ENSICAEN, CNRS, Normandie
Univ, GREYC UMR 6072, F-14000 Caen, France

Abstract

Floating-point arithmetic plays a central role in computer science and is used in various domains where precision and computational scale are essential. One notable application is in machine learning, where Fully Homomorphic Encryption (FHE) can play a crucial role in safeguarding user privacy. In this paper, we focus on TFHE and develop novel homomorphic operators designed to enable the construction of precise and adaptable homomorphic floating-point operations. Integrating floating-point arithmetic within the context of FHE is particularly challenging due to constraints such as small message space and the lack of information during computation. Despite these challenges, we were able to determine parameters for common precisions (e.g., 32-bit, 64-bit) and achieve remarkable computational speeds, with 32-bit floating-point additions completing in 2.5 seconds and multiplications in approximately 1 second in a multi-threaded environment. These metrics provide empirical evidence of the efficiency and practicality of our proposed methods, which significantly outperform previous efforts. Our results demonstrate a significant advancement in the practical application of FHE, making it more viable for real-world scenarios and bridging the gap between theoretical encryption techniques and practical usability.

^{*}Ilaria Chillotti and Damien Ligier contributed to this work during a previous employment at Zama.

Contents

1	Introduction	3
1.1	Our Results	3
1.1.1	Prior Approaches	6
1.1.2	Roadmap	8
2	Background and Notations	9
2.1	Notations	9
2.2	FHE Ciphertext Types	9
2.3	FHE Operators	10
2.4	Representing Large Integers with TFHE	13
2.4.1	Integer Subtraction	15
2.5	Traditional Floating-Point Representation	18
3	Homomorphic Floating-Points (HFP)	19
3.1	MiniFloats: WoP-PBS Based Floats	19
3.2	Homomorphic Floating-Point Encoding	20
3.3	Choosing Between Two Ciphertexts	24
3.4	Propagating the Carries	25
4	Addition and Subtraction of HFP	29
4.1	Managing Mantissas and Exponents	29
4.1.1	Aligned Mantissa	29
4.1.2	SubMantissa	33
4.2	Addition and Subtraction	37
5	Multiplication and Division	40
5.1	Multiplication	41
5.2	Division	44
6	Experimental Results	46
7	More Features over HFP	49
7.1	Managing Special Values	49
7.2	Computing Function Approximations	50
7.3	Other Operations	50
7.3.1	ReLU	51
7.3.2	Approximate Sigmoid	52

1 Introduction

Fully homomorphic encryption (FHE) is a cryptographic paradigm that allows evaluating any circuit over encrypted data, ensuring that the decrypted output contains the results of the circuit evaluation. The concept of FHE was first proven possible in the seminal paper by Gentry [Gen09]. Since then, substantial work has been done to enhance performance. As of today, the following schemes are considered to be practical, BGV [BGV12], BFV [Bra12, FV12], CKKS [CKKS17, CHK⁺18], FHEW [DM15], TFHE [CGGI20]. All these schemes are based on the Learning With Errors (LWE) problem [Reg05], its Ring variants (RLWE) [SSTX09, LPR10] and the General LWE (GLWE), also called Module LWE [LS15, BGV12]. However, there is a variation in their approach during computations: while BGV/BFV/FHEW/TFHE perform exact integer computations, CKKS generally focuses on approximate computations using fixed-point arithmetic. The easiest solution to adapt the other schemes to support real numbers is to define a special encoding mapping modular integers to fixed-point numbers as done in [CJP21, CJL⁺20] for TFHE, or in [CSVW16] for BGV and [Lai17] for BFV. A significant limitation of this method, inherent to fixed-point arithmetic, is that the encoding range must be sufficiently large to cover all values used in the application. Otherwise, the cumulative precision loss during computations becomes excessively large. In the case of TFHE, this problem is twofold: one must either extend the precision and/or bound the circuit depth. The former significantly impacts performance, as the bootstrapping efficiency is closely linked to the number of bits required to represent a message. In practical terms, TFHE's bootstrapping is deemed efficient when the precision is smaller than 10 bits [BBB⁺22]. The latter is a direct contradiction with the TFHE construction, whose main advantage is to have the possibility to evaluate any circuit depths through the functional/programmable bootstrapping. Therefore, the only feasible approach to handle real numbers within TFHE is to employ floating-point representation. Thus, the problem we are solving in this paper is: *How to design an efficient and customizable floating-point arithmetic based on TFHE?*

1.1 Our Results

In the current landscape, FHE schemes, and especially TFHE, are not designed to seamlessly integrate with established floating-point standards,

which are tailored to match hardware specifications, such as bit sizes and the frequent use of conditional instructions. This paper introduces novel homomorphic operators that can serve as the building blocks for constructing homomorphic floating-point arithmetic with arbitrary precision within a given FHE framework. As a matter of fact, our method is compliant with any FHE scheme providing a bootstrapping which is exact and programmable, so FHEW/TFHE and any of their more optimized versions ([BIP⁺22, LMK⁺23]) should work.

Our innovative operators not only align with widely recognized floating-point arithmetic standards like IEEE754 [Ins08], but also adapt to variants with custom precision levels, boosting performance. This marks the first feasible approach for performing homomorphic computation with floating-points, significantly surpassing previous methods in time efficiency by a factor of at least 10.

Currently, there is a substantial gap between cleartext floating-point arithmetic and homomorphic floating-point arithmetic. Even if there remains a considerable distance to achieving a fully practical deployment, our work consistently narrows this gap. While not yet practical for general-purpose applications, our solution may already be sufficiently efficient for certain specialized use cases.

In TFHE, the process to encrypt large integers involves dividing the integer message into several parts, each of which is then encrypted as its own LWE ciphertext. This method is crucial for handling large numerical values effectively by encrypting them in smaller, discrete units [BST20, GBA21, KO22, CZB⁺22, LMP21, CLOT21, BBB⁺22]. Then, the operations are either modular, or the number of blocks increases all along the computation. To ensure correctness, part of the cleartext space, referred to as the carry space, is reserved to accommodate the growth of the message during computation. Each ciphertext is also paired with public metadata, referred to as the *degree* in previous studies [BBB⁺22], which tracks the maximum potential size of the message. When the public metadata indicates that the carry space is nearing capacity, the carry buffer is either cleared or a new block is added at each operation. This approach is based on worst-case size estimations, meaning that in many cases, additional operations could be performed before needing to clear the carry space or add a new block.

To construct floating-point arithmetic, we must represent mantissas and exponents, which can be viewed as large integers. We aim to make this arithmetic as efficient as possible, which requires circumventing the limitations of

worst-case estimations. In fact, an inherent challenge of FHE is the lack of information about underlying values, due to the data being encrypted. This limitation prevents the typical optimizations commonly used in floating-point arithmetic implementations, such as those described in [PFH⁺20].

We introduce novel algorithms that automatically retain only the most significant bits and discard some of the least significant bits, maintaining the same representation throughout. This mimics a well-known round mode of floats, called the rounding towards zero mode. We refer to [MBDD⁺18, Section 2.2.1] or to [Hwa24] for more details. Our first method, based on [BBB⁺22], ensures the correct encoding after each operation. This method is effective for small mantissas and exponents but does not scale well. The second method leverages circuit bootstrapping to obtain the carry value, allowing the ciphertext to be updated accordingly. This approach is faster than the first method for large mantissas or exponents.

Building on large integers, we introduce two different methods to build efficient floating-point arithmetic, each with its own pros and cons. The first approach, detailed in Section 3.1, heavily utilizes the alternative PBS proposed in [BBB⁺22] (referred to as WoP-PBS) to perform operations on floating-point numbers. This method allows for the efficient evaluation of functions on ciphertexts that encrypt large integers. Our approach, while straightforward, distinguishes itself from other homomorphic floating-point methods by representing a floating-point number within a single ciphertext (or more). This technique proves particularly efficient for floating-point numbers with small precision (up to 12 bits).

The second approach, which is the core of this contribution, constructs floating-point arithmetic based on TFHE that can follow standards such as [Ins08] and is not constrained by precision. Our new algorithms make extensive use of an extended version of CMux, a homomorphic operator that selects between two inputs based on an encrypted decision bit. This method is key to developing faster homomorphic operations that effectively combine the sign, mantissa, and exponent of one or more homomorphic floats. During homomorphic float operations, to compute the resulting mantissa, an extra LWE ciphertext will be used to make the operation both faster and more precise. This approach, soberly titled homomorphic floating-point (HFP), is the first that can be considered to be deployed for real use-cases. Practically, on a typical server machine, an addition of two 32-bit HFP numbers takes around 2 seconds, while a multiplication takes around 1 second. Beyond

the description of all arithmetic operations, including the division, we also provide their noise analysis and the hypotheses done to generate the cryptographic parameters and their associated failure probability. We also include algorithms to efficiently compute the ReLU and the sigmoid, two of the most used functions in machine learning. We show how to easily extend our approach to take into account some floating-point subtleties, like the special values. As a simple application and showcase of the versatility of floating-points, we briefly detail how to compute the approximation of any functions. In the end, our method outperforms the state-of-the-art (as shown in the next section), which was more about showing that floating-points might be doable with TFHE, rather than giving a practical solution as we do here.

1.1.1 Prior Approaches

Efforts to develop efficient FHE computation methods for real numbers can be categorized into two approaches: the fixed-point arithmetic approach and the floating-point arithmetic approach. Most of the first attempts [CSVW16, Lai17, AN16] focus on the BFV scheme [Bra12, FV12]. In [CSVW16, Lai17], the authors chose the fixed-point approach. Roughly, their idea is to decompose a real number into two integers, one representing the value before the point and the second representing the value after the point. The binary decomposition of the two integers is encrypted in one RLWE ciphertext such that the integer part is encrypted over the coefficient of small degree and the fractional part is encrypted on the coefficient of high degree. This method encounters two significant limitations: first, after several operations, accuracy is compromised due to the mixing of the fractional and the integer parts of the number. Second, the computation must remain within a certain modulus limit; exceeding this threshold also results in a loss of correctness. Thus, using fixed-point arithmetic is particularly suitable for FHE schemes where the depth of the circuit is somewhat bounded, since they share a similar constraint. In fact, an encrypted floating-point number is often represented by one ciphertext for the sign, one or several ciphertexts for the mantissa and one or several ciphertexts for the exponent. Their approach is also based on FV, whose bootstrapping is neither efficient nor programmable. This results in unpractical methods that cannot be adapted to TFHE.

In the TFHE context, to the best of our knowledge, only two techniques have been studied [ML20] and [LS22]. The former [ML20] uses the tradi-

tional floating-point representation, where each LWE ciphertext contains one boolean value. Then, they rely only on boolean gate operations to perform the floating-point computation. Beyond the lack of space efficiency, the major problem with this solution is its time complexity. In the latter [LS22], the authors take advantage of RLWE ciphertexts to represent floats. They have a floating-point in three parts: an RLWE for the sign, one for the mantissa and another one for the exponent. In this representation, the sign and the exponent are each represented on the first coefficient of their RLWE. The mantissa is represented by several coefficients depending on the base of the decomposition. In this work, they propose a method to detect overflow based on an encrypted witness. As previously, the main drawback lies in the time complexity, which suffers from an exponential factor related to the size of the exponent and the use of tensor product and relinearisation.

In our work, we use the standard representation proposed before (several LWE ciphertexts for the mantissa, several LWE ciphertexts for the exponent and one LWE ciphertext for the sign). The main change in our algorithm is the use of circuit bootstrapping (CBS) which is costlier than a PBS ($\mathbb{C}_{\text{CBS}} \approx \ell_{\text{CBS}} \cdot \mathbb{C}_{\text{PBS}}$) but gives the possibility to perform CMuxes and reduce the cost of the TFHE algorithm used in the floating-point arithmetic. Note that our work could benefit from recent optimizations done in [WWL⁺24].

The code for [ML20] is not publicly available, and despite significant effort, we were unable to successfully run the code provided with [LS22]. As a result, the comparison below is based on the timings reported in the respective papers. In [ML20], the experiments were conducted on an Intel i7-6700@3.40 GHz (up to 4 GHz) with 8 threads. Since the computational model (i.e., sequential or parallel) is not specified, we assume these experiments were run sequentially. In the latter paper [LS22], experiments were run on an Intel Xeon Silver 4210@2.40 GHz, with 40 threads. As shown in Table 1, our approach significantly outperforms existing methods, achieving at least an 8-fold improvement (for 32-bit floating-point multiplication) and up to 100-fold improvement (for 64-bit floating-point addition). Sequential timings are:

As explained above, each technique was evaluated on different machines, with some assumptions regarding the nature of the computation (sequential or parallel). To provide a more complete comparison, Table 2 presents a complexity comparison in terms of the number of PBS operations required for the main operations (addition and multiplication) between our work and the state of the art. This comparison highlights the removal of the exponential

Paper (Precision)	Add	Mul
[ML20] (32-bits)	490s	162s
[LS22] (32-bits)	530s	443s
Ours (32-bits)	7s	20s
[ML20] (64-bits)	1200s	686s
[LS22] (64-bits)	858s	808s
Ours (64-bits)	12s	82s

Table 1: Comparison of addition and multiplication times with the state-of-the-art.

factor in the complexity formulas, further emphasizing the efficiency of our approach.

In this context, the terms $\rho_m \cdot \ell_m$ and $\rho_e \cdot \ell_e$ represent the number of bits in the mantissa and exponent, respectively. In our method, ρ_m (and ρ_e) corresponds to the message space in each LWE ciphertext for the mantissa (and exponent), while ℓ_m (and ℓ_e) indicates the number of ciphertexts representing the mantissa (and exponent). Considering the complexity of the two previous works based on TFHE, we achieve a significant gain by eliminating the exponential factor for addition (present in both previous techniques) and for multiplication specifically in [LS22]

	Addition Complexity	Multiplication Complexity
[ML20]	$\approx 2^{\rho_e \ell_e} (\ell_m \rho_m + \ell_e \rho_e) \log(\ell_m \rho_m + \ell_e \rho_e) \mathbb{C}_{\text{PBS}}$	$\approx (\ell_m \rho_m)^2 + \ell_e \rho_e \log(\ell_e \rho_e) \mathbb{C}_{\text{PBS}}$
[LS22]	$> 2^{\rho_e \ell_e + 1} \mathbb{C}_{\text{PBS}}$	$> 2^{\rho_e \ell_e + 1} \mathbb{C}_{\text{PBS}}$
Ours	$(3 \cdot \ell_m + 6 \cdot \ell_e + 3) \mathbb{C}_{\text{PBS}} + (\ell_m + \ell_e \cdot \rho_e + 3) \mathbb{C}_{\text{CBS}}$	$\left(\frac{\ell_m^2}{2} \left(1 + \frac{1}{\rho_e} \right) + \ell_m \left(2 + \frac{1}{\rho_e} \right) + 4 \right) \mathbb{C}_{\text{PBS}} + \mathbb{C}_{\text{CBS}}$

Table 2: Our Method vs. Existing Works.

1.1.2 Roadmap

In Section 2, we give an overview of TFHE and floating-point arithmetic. Section 3 presents our initial efforts in developing efficient floating-point arithmetic using WoP-PBS from [BBB⁺22], followed by a new, more versatile encoding methodology for floating-points tailored for TFHE, along with fundamental building blocks for advanced operations. Sections 4 and 5 focus on detailing algorithms for floating-point number computation. Section 6 highlights the practicality of our methods, demonstrating their efficiency with

timing benchmarks for standard float precision. Finally, Section 7 describes how to improve our representation with the infinity values and how to compute approximate functions.

2 Background and Notations

2.1 Notations

Let q be a positive integer. We note \mathbb{Z}_q the ring $\mathbb{Z}/q\mathbb{Z}$. Let N be a power of two, representing the degree of quotient polynomial. Then, we note $\mathfrak{R}_{q,N}$ the ring $\mathbb{Z}_q[X]/(X^N + 1)$. By convention, the vectors are written in bold \mathbf{x} , polynomials in upper case X , and integers are in lower case x . Moreover, the values associated with floating-points will be written in the following manner: \mathbf{m} for the mantissa, \mathbf{e} for the exponent, \mathbf{s} for the sign, and \mathbf{z} for any other value. The number of ciphertexts associated with \mathbf{z} is denoted by $l_{\mathbf{z}}$, and the number of bits encoding this value by $\rho_{\mathbf{z}}$. Regarding probability distributions, the uniform distribution over a set S is denoted as $\mathcal{U}(S)$ whereas a Gaussian distribution with a mean set to zero and a variance σ^2 is written \mathcal{N}_{σ^2} .

2.2 FHE Ciphertext Types

TFHE security is based on the Learning With Errors (LWE) assumption [Reg05], its extension to polynomial rings [SSTX09, LPR10] RLWE and the Generalized approach GLWE [BGV12, LS15].

Definition 1 (GLWE Ciphertexts [BGV12, LS15]) Let $\mathbf{S} = (S_0, \dots, S_{k-1}) \in \mathfrak{R}_{q,N}^k$ be the secret key, with $S_i = \sum_{j=0}^{N-1} s_{i,j} X^j$, where each coefficient $s_{i,j}$ is sampled from a uniform binary, uniform ternary or Gaussian distribution. Let $\mathbf{A} = (A_0, \dots, A_{k-1}) \leftarrow \mathcal{U}(\mathfrak{R}_{q,N})^k$ be the mask and let $E \in \mathfrak{R}_{q,N}$ be the noise, where each coefficient e_i is sampled from a Gaussian distribution \mathcal{N}_{σ^2} . Let $\Delta \in \mathbb{N}$ be the scaling factor depending on the plaintext space p , such that $\Delta = \frac{q}{2p}$. A GLWE ciphertext of a plaintext $\Delta \cdot M \in \mathfrak{R}_{q,N}$ under a secret key $\mathbf{S} \in \mathfrak{R}_{q,N}^k$ is defined as:

$$\text{CT} = \left(\mathbf{A}, B = \sum_{i=0}^{k-1} A_i \cdot S_i + \Delta \cdot M + E \right) \in \text{GLWE}_{\mathbf{S}}(\Delta \cdot M) \subseteq \mathfrak{R}_{q,N}^{k+1}.$$

In what follows, Δ is implicit.

Definition 2 (GLEV ciphertexts [CLOT21]) For a given decomposition base $\beta \in \mathbb{N}^*$ and a level decomposition $\ell \in \mathbb{N}^*$, a GLEV ciphertext of a message $M \in \mathfrak{R}_{q,N}$ under a secret key $\mathbf{S} \in \mathfrak{R}_{q,N}^k$ is a ciphertext composed of ℓ GLWE ciphertexts encrypting the same message M for different scaling factors (given by the base β and the level ℓ). Let $\text{CT}_i \in \text{GLWE}_{\mathbf{S}}\left(\frac{q}{\beta^{i+1}}M\right) \subseteq \mathfrak{R}_{q,N}^{k+1}$ for $i \in [0, \ell)$. Then, a GLEV ciphertext is defined as:

$$\overline{\text{CT}} = (\text{CT}_0, \dots, \text{CT}_{\ell-1}) \in \text{GLEV}_{\mathbf{S}}^{\beta, \ell}(M) \subseteq \mathfrak{R}_{q,N}^{\ell \times (k+1)}.$$

Definition 3 (GGSW ciphertexts [GSW13, CLOT21]) For a given decomposition base $\beta \in \mathbb{N}^*$ and a level decomposition $\ell \in \mathbb{N}^*$, a GGSW ciphertext encrypting a message $M \in \mathfrak{R}_{q,N}$ under a secret key $\mathbf{S} \in \mathfrak{R}_{q,N}^k$ is composed of $(k+1)$ GLEV ciphertexts encrypting the same message M multiplied by elements of the secret key for different scaling factors (given by the base β and the level ℓ). Let $\overline{\text{CT}}_j \in \text{GLEV}_{\mathbf{S}}^{\beta, \ell}(-S_j \cdot M) \subseteq \mathfrak{R}_{q,N}^{\ell \times (k+1)}$ for $j \in [0, k)$ and $\overline{\text{CT}}_k \in \text{GLEV}_{\mathbf{S}}^{\beta, \ell}(M) \subseteq \mathfrak{R}_{q,N}^{\ell \times (k+1)}$. Then, a GGSW ciphertext is defined as:

$$\overline{\overline{\text{CT}}} = (\overline{\text{CT}}_0, \dots, \overline{\text{CT}}_k) \in \text{GGSW}_{\mathbf{S}}^{\beta, \ell}(M) \subseteq \mathfrak{R}_{q,N}^{(k+1)\ell \times (k+1)}.$$

Remark 1 (LWE and RLWE) A GLWE ciphertext with $N = 1$ is a LWE ciphertext. In this case, we consider $n = k$ for the LWE size, and all the elements of the ciphertext are denoted in lower case (i.e., $\text{ct}, \mathbf{s}, \mathbf{a}$ and \mathbf{b}). A GLWE ciphertext with $k = 1$ and $N > 1$ is a RLWE ciphertext. This can be extended to GLEV and GGSW ciphertext as well.

2.3 FHE Operators

In what follows, we recall some of the TFHE operators. We refer the reader to the associated references for more details. We assume that all the correctness conditions are fulfilled when one of these algorithms is called. More details on these conditions are given in lemmas alongside almost every algorithm introduced in this paper (see Lemma 6).

Extract and Insert LWE samples The sample extract is an algorithm taking as input a GLWE ciphertext in $\mathfrak{R}_{q,N}^{k+1}$ and returning its i^{th} coefficient

as an LWE ciphertext in \mathbb{Z}_q^{kN+1} for i in $[0, N)$. This operation is noiseless and is detailed in [CGGI20].

$$\text{ct}_{\text{out}} \leftarrow \text{SampleExtract}(\text{CT}_{\text{in}}, i).$$

Usually, only the first coefficient is extracted. In this case, this is simply denoted: $\text{SampleExtract}(\text{CT}_{\text{in}}) = \text{SampleExtract}(\text{CT}_{\text{in}}, 0)$.

Conversely, inserting a sample is an operation that creates a GLWE ciphertext from a LWE ciphertext taken as input. The first coefficient of the GLWE ciphertext encrypts the same plaintext as the LWE ciphertext. All the other coefficients encrypt random values. See [BCL⁺23, Alg. 9] for more details. The signature of this algorithm is:

$$\text{CT}_{\text{out}} \leftarrow \text{SampleInsert}(\text{ct}_{\text{in}}).$$

Keyswitch (KS) The keyswitch (e.g., detailed in [CGGI20, CLOT21]) is an homomorphic operation which changes the encryption secret key of a ciphertext using only public material. This operation takes as input an LWE ciphertext encrypting a message m under a secret key $\mathbf{s}_{\text{in}} \in \mathbb{Z}_q^{n_{\text{in}}}$ and returns an LWE ciphertext encrypting the same message m under a secret key $\mathbf{s}_{\text{out}} \in \mathbb{Z}_q^{n_{\text{out}}}$. To perform this operation, a key-switching-key (KSK) is required. A KSK is an encryption of \mathbf{s}_{in} with redundancy under the secret key \mathbf{s}_{out} , i.e., $\text{KSK} = (\text{KSK}_0, \dots, \text{KSK}_{n_{\text{in}}-1})$ where $\text{KSK}_i \in \text{LEV}_{\mathbf{s}_{\text{out}}}^{\beta, \ell}(\mathbf{s}_{\text{in}, i})$ for $i \in [0, n_{\text{in}})$. The signature is:

$$\text{ct}_{\text{out}} \leftarrow \text{KS}(\text{ct}_{\text{in}}, \text{KSK}).$$

Programmable Bootstrapping (PBS) Unique among bootstrapping techniques, TFHE's bootstrapping not only reduces the noise in a ciphertext but also has the distinct capability to evaluate any univariate function f represented as a Look-Up Table LUT_f . This distinctive feature of evaluating a function is the reason TFHE's bootstrapping is often referred to as the functional or programmable bootstrapping (PBS) [CGGI20, CJL⁺20, CJP21]. The PBS is done in three steps: a modulus switch (MS), a blind rotation (BR) and a sample extract (SE). Taking as input the bootstrapping key BSK, the look-up table LUT_f and an encrypted LWE ciphertext of a message m under a secret key \mathbf{s}_{in} , the bootstrapping outputs an LWE ciphertext which encrypts the message $f(m)$ under the secret key \mathbf{s}_{out} with a smaller noise.

$$\text{ct}_{\text{out}} \leftarrow \text{PBS}(\text{ct}_{\text{in}}, \text{LUT}_f, \text{BSK}).$$

The bootstrapping key is an encryption of each element of the input secret key $\mathbf{s}_{\text{in}} = (s_0, \dots, s_{n-1})$ as GGSW ciphertexts encrypted under a key $\mathbf{S}_{\text{out}} \in \mathfrak{R}_{q,N}^k$ (note that each element of the input secret key \mathbf{s}_{in} is seen as an element of $\mathfrak{R}_{q,N}$ where s_i corresponds to the constant coefficient, and all the other coefficients are set to zero). So we have $\text{BSK} = (\text{BSK}_0, \dots, \text{BSK}_{n-1})$ where $\text{BSK}_i \in \text{GGSW}_{\mathbf{S}}^{\beta,\ell}(s_i)$ for $i \in [0, n)$ and where s_i corresponds to an element of the LWE input secret key \mathbf{s} . As introduced in [CJP21], to improve the computational time of the PBS, a keyswitch always precedes a PBS. In the following, the two operations are merged together.

$$\text{ct}_{\text{out}} \leftarrow \text{KS-PBS}(\text{ct}_{\text{in}}, \text{LUT}_f, \text{KSK}, \text{BSK}).$$

Circuit Bootstrapping (CBS) [CGGI20, Alg. 11] The circuit bootstrapping (CBS) is an operation that takes as input an LWE ciphertext encrypting a binary message $m \in \{0, 1\}$ under a secret key \mathbf{s}_{in} and returns a GGSW ciphertext encrypting the same binary message m under the secret key \mathbf{S}_{out} .

$$\overline{\overline{\text{CT}}}_{\text{out}} \leftarrow \text{CBS}(\text{ct}_{\text{in}}, \text{BSK}, \text{KSK}, \text{PFKSK}).$$

BSK is the bootstrapping key as defined above and PFKSK stands for private functional key switch key. A PFKSK = $[\text{PFKSK}_0, \dots, \text{PFKSK}_k]$ is a public key encrypting the product of each input and output secret key where $\text{PFKSK}_i = \text{PFKSK}_{\mathbf{s}_{\text{in}} \rightarrow \mathbf{S}_{\text{out}}}(S_{\text{out},i}) = [\text{GLEV}_{\mathbf{S}_{\text{out}}}^{\beta,\ell}(-s_{\text{in},0} \cdot S_{\text{out},i}), \dots, \text{GLEV}_{\mathbf{S}_{\text{out}}}^{\beta,\ell}(-s_{\text{in},kN-1} \cdot S_{\text{out},i})]$ for $i \in [0, k_{\text{out}}]$ with $S_{\text{out},k} = -1$. With this public key, we can perform an operation with an LWE ciphertext encrypting a message m under the secret key \mathbf{s} and PFKSK_i to obtain a GLWE ciphertext under the secret key \mathbf{S}_{out} of the message $m \cdot S_{\text{out},i}$. The CBS operation is done in two steps. The first one consists in performing several bootstrappings on the input LWE ciphertext to obtain many LWE ciphertexts which encrypt the binary message with different scaling factors depending on the decomposition base β and the level decomposition ℓ of the GGSW ciphertext. This operation transforms the input LWE ciphertext $\text{ct} \in \text{LWE}_{\mathbf{s}}(m)$ into a LEV ciphertext $\overline{\text{ct}} \in \text{LEV}_{\mathbf{s}}^{\beta,\ell}(m)$. The second step consists in using the PFKSK to multiply each of the LWE ciphertext of the LEV obtained after the first operation to get $k+1$ new GLEV ciphertexts $\overline{\overline{\text{CT}}}$ encrypting $m \cdot S_{\text{out},i}$ for $i \in [0, k]$ (with $S_{\text{out},k} = -1$). This collection of $\overline{\overline{\text{CT}}}$ corresponds to a GGSW ciphertext as defined in Definition 3.

Remark 2 In what follows, we constrain the input LWE secret key $\mathbf{s}_{\text{in}} = (s_0, \dots, s_{kN-1})$ to contain the same coefficients as the flattened GGSW output

secret key $\mathbf{S}_{\text{out}} = (S_0, \dots, S_{k-1})$ where $S_i = \sum_{j=0}^{N-1} s_{iN+j} X^j$. Thanks to this constraint, we can use the sample insert operation to cast an LWE ciphertext encrypted under \mathbf{s}_{in} as a GLWE ciphertext encrypted under \mathbf{S}_{out} .

CMux The CMux selects one of two GLWE ciphertexts, depending on the value of an encrypted decision bit b . Let $\overline{\overline{\text{CT}}} \in \text{GGSW}_{\mathbf{S}}(b)$ with $b \in \{0, 1\}$, $\text{CT}_0 \in \text{GLWE}_{\mathbf{S}}(\Delta \cdot m_0)$ and $\text{CT}_1 \in \text{GLWE}_{\mathbf{S}}(\Delta \cdot m_1)$. The output ciphertext CT_{out} is an encryption of $\Delta \cdot m_b$. The signature of the algorithm is:

$$\text{CT}_{\text{out}} \leftarrow \text{CMux}(\text{CT}_0, \text{CT}_1, \overline{\overline{\text{CT}}}).$$

Remark 3 (Public Keys) In what follows, the public key PUB encompasses the bootstrapping key (BSK), the keyswitching key (KSK) and the private functional key switch key (PFKSK). To simplify the algorithms, we will simply refer to PUB when the context is clear enough to decide which key should be used.

2.4 Representing Large Integers with TFHE

Only integer plaintexts smaller than 10 bits can be encoded in TFHE [BBB⁺22]. This is due to the bootstrapping, where the plaintext precision is entangled with the degree N of the cyclotomic polynomial in $\mathfrak{R}_{q,N}^k$. We briefly recall the techniques described in [BBB⁺22] to overcome this limitation. We first describe the encoding of large integers, and then focus on the programmable bootstrapping associated with such encoding.

Large Integer Encoding and Arithmetic Let $m \in \mathbb{Z}_p$ be the cleartext to encrypt such that $\log_2(p) > 10$. The idea is to apply a radix decomposition to m before encrypting each part into a dedicated LWE ciphertext. Let $\rho \in \mathbb{N}$ such that $2 \leq \rho \leq p < q$. Intuitively, ρ defines the message space, whereas $\frac{p}{2\rho}$ refers to the carry space, which is getting used all along homomorphic computation. Let $\ell_z = \lceil \log_\rho(m) \rceil$. The encoded plaintext \mathbf{pt}_m is defined as the ρ -radix decomposition of a message $m \in \mathbb{N}$, i.e., $\mathbf{pt}_m = [m_{\ell_z-1} \dots m_0] \leftarrow \text{Encoding}^\rho(m)$ with $m = \sum_{i=0}^{\ell_z-1} m_i \rho^i$. Then, each m_i is independently encrypted, so that $\mathbf{ct}_m = [\text{ct}_{m,\ell_z-1}, \dots, \text{ct}_{m,0}] \in [\text{LWE}_{\mathbf{s}}(m_{\ell_z-1}), \dots, \text{LWE}_{\mathbf{s}}(m_0)]$. Common modular integer operations are then defined on these ciphertexts,

in particular modular additions, subtractions, multiplications and division are possible.

By construction, after a given number of operations, the carry space is full and we need to call the **CarryPropagate** algorithm to homomorphically propagate the carries. This algorithm was first introduced in [BBB⁺22]. The goal of this algorithm is to clear the carry space of each input ciphertext without losing information, except for the most significant carry, which is lost. This loss occurs because, in [BBB⁺22], the authors work with integers modulo $\rho_3^{\ell_3}$.

When dealing with floating-point numbers, however, losing this information is not acceptable, especially when the mantissa is full. By retaining this information, we can grow the exponent if necessary. To address this limitation, we modified the **CarryPropagate** algorithm from [BBB⁺22] and named it **CarryPropagateExtended**, which is detailed in Algorithm 1. Note that to recover the original **CarryPropagate** algorithm, it is sufficient to execute the loop from 0 to $\ell_3 - 2$.

In what follows, we assume $p = \rho^2$ to simplify the carry propagation. We will refer to the algorithm computing an integer operation **Op** as **IntOp**.

Programmable Bootstrapping over Large Integers (WoP-PBS)

In [BBB⁺22], the authors present an algorithm designed for performing function evaluation on large integers and reducing the noise. The Without Padding Programmable Bootstrapping (WoP-PBS) algorithm, as described in [BBB⁺22], is named for its unique feature of not requiring a padding bit which distinguishes it from the typical TFHE's PBS. Having no padding bit enables the use of the total plaintext space instead of half of it, i.e., the scaling factor Δ becomes $\frac{q}{p}$ instead of $\frac{q}{2p}$ as described in Definition 1. Briefly, given a plaintext m represented using ρ bits, a WoP-PBS starts by extracting each bit of the plaintext m to obtain a list of LWEs where each one of them encrypts a bit, i.e., $\{\text{ct}_i\}_{i \in [0, \dots, \rho-1]} \in \text{LWE}_s(m_i) \subset \mathbb{Z}_q^{(n+1) \times \rho}$. Then a circuit bootstrapping is applied to convert them into GGSW ciphertexts. The next step consists in computing a vertical packing (described in [CGGI20, Alg. 5]) to choose the right LUT depending on the value of the message. This ends by a classical blind rotation to select the correct value into the LUT. The detailed algorithm can be found in [BBB⁺22, Alg. 3]. The signature of the WoP-PBS is:

$$\text{ct}_{\text{out}} \leftarrow \mathbf{WoP-PBS}(\text{ct}_{\text{in}}, \text{LUT}_f, \text{BSK}, \text{KSK}, \text{PFKSK}) = \mathbf{WoP-PBS}(\text{ct}_{\text{in}}, \text{LUT}_f, \text{PUB}).$$

Algorithm 1: $\mathbf{ct}_{\text{out}} \leftarrow \text{CarryPropagateExtended}(\mathbf{ct}_{\text{in}}, \text{PUB})$

Context: $\begin{cases} \Delta_{\mathfrak{z}} : & \text{scaling factor} \\ \text{LUT}_{\text{Carry}} : & \text{LUT to return the carry of the ciphertext (return } \lfloor \frac{m}{\beta_{\mathfrak{z}}} \rfloor \text{)}. \\ \text{LUT}_{\text{Msg}} : & \text{LUT to return the message of the ciphertext (return } m \bmod p_{\mathfrak{z}} \text{)}. \\ p_{\mathfrak{z}} : & \text{the carry-message modulus} \\ \beta_{\mathfrak{z}} : & \text{the message modulus} \end{cases}$

Input: $\begin{cases} \mathbf{ct}_{\text{in}} = [\mathbf{ct}_{\text{in}, \ell_{\mathfrak{z}}-1}, \dots, \mathbf{ct}_{\text{in}, 0}] \in \mathbb{Z}_q^{(n+1) \cdot \ell_{\mathfrak{z}}} \\ \text{PUB} : \text{Public materials for PBS-KS} \end{cases}$

Output: $\begin{cases} \mathbf{ct}_{\text{out}} = [\mathbf{ct}_{\text{tmp}}, \mathbf{ct}_{\text{out}, \ell_{\mathfrak{z}}-1}, \dots, \mathbf{ct}_{\text{out}, 0}] \in \mathbb{Z}_q^{(n+1) \cdot (\ell_{\mathfrak{z}}+1)} \end{cases}$

```

1 for  $i$  in  $[0.. \ell_{\mathfrak{z}} - 1]$  do
    /* Extract the carry */
2    $\mathbf{ct}_{\text{tmp}} \leftarrow \text{KS-PBS}(\mathbf{ct}_{\text{in}, i}, \text{PUB}, \text{LUT}_{\text{Carry}})$ 
    /* Extract the message */
3    $\mathbf{ct}_{\text{out}, i} \leftarrow \text{KS-PBS}(\mathbf{ct}_{\text{in}, i}, \text{PUB}, \text{LUT}_{\text{Msg}})$ 
4   if  $i \neq \ell_{\mathfrak{z}}-1$  then
5      $\mathbf{ct}_{\text{in}, i+1} \leftarrow \mathbf{ct}_{\text{in}, i+1} + \mathbf{ct}_{\text{tmp}}$ 
6 return  $(\mathbf{ct}_{\text{out}} = [\mathbf{ct}_{\text{tmp}}, \mathbf{ct}_{\text{out}, \ell_{\mathfrak{z}}-1}, \dots, \mathbf{ct}_{\text{out}, 0}])$ 

```

2.4.1 Integer Subtraction

In the following floating-point algorithms, we require an operation that takes as input two vectors of ciphertexts and returns the sign of the subtraction along with a vector of ciphertexts representing the absolute value of the subtraction. This method is detailed in Algorithm 2. Intuitively, in the initial steps, an offset is added to ensure that the messages in each ciphertext of the first input are larger than those of the second input. Next, we perform the subtraction between the adjusted first input and the second input, followed by a carry propagation as described in Algorithm 1. We then extract the most significant bit (MSB) of the top block from the resulting ciphertext. Finally, we traverse all the blocks, returning the value if the MSB is set to 1, or the opposite if it is not. The sign is encoded in the padding bit, with an offset such that the most significant bit is 0 for positive values and 1 for negative values.

Lemma 1 (IntSub* (Algorithm 2)) Let $\mathbf{ct}_{\mathfrak{z}_i} = [\mathbf{ct}_{\mathfrak{z}_i, \ell_3-1}, \dots, \mathbf{ct}_{\mathfrak{z}_i, 0}] \in [\text{LWE}_s(\mathfrak{z}_{i, \ell_3-1}), \dots, \text{LWE}_s(\mathfrak{z}_{i, 0})] \subseteq \mathbb{Z}_q^{(n+1) \cdot \ell_3}$ with $i \in \{0, 1\}$ be two ciphertexts encrypting $\mathfrak{z}_i \in \mathbb{N}$.

Let $(\mathbf{ct}_{\mathfrak{z}_{\text{sub}}}, \mathbf{ct}_s) \leftarrow \text{IntSub}^*(\mathbf{ct}_{\mathfrak{z}_1}, \mathbf{ct}_{\mathfrak{z}_2}, \text{PUB})$. Then $\text{Decrypt}_s(\mathbf{ct}_{\mathfrak{z}_{\text{sub}}}) = |\mathfrak{z}_1 - \mathfrak{z}_2|$ and $\text{Decrypt}_s(\mathbf{ct}_s) = \text{Sign}(\mathfrak{z}_1 - \mathfrak{z}_2)$.

The complexity of Algorithm 2 can be defined as $\mathbb{C}_{\text{IntSub}^*}^{\mathfrak{z}} = (3 \cdot \ell_3 + 1) \cdot \mathbb{C}_{\text{PBS}}$

Proof 1 (Correctness of IntSub* (Algorithm 2)) In this algorithm, $\text{CT}_{\mathfrak{z}}$ can represent the mantissa or the exponent. The goal of this algorithm is to return $|\mathfrak{z}_1 - \mathfrak{z}_2|$ and $\text{Sign}(\mathfrak{z}_1 - \mathfrak{z}_2)$. The first step of this algorithm is to ensure that the message \mathfrak{z}_1 is bigger than \mathfrak{z}_2 . To do that, we add $2^{(\ell_3+1) \cdot \rho_3-1}$ to \mathfrak{z}_1 , which automatically guarantees that all the LWE that compose \mathfrak{z}_1 such that all the LWE of \mathfrak{z}_1 are bigger than the ones of \mathfrak{z}_2 . Now we can perform the subtraction term by term. After a carry propagate, if $\mathfrak{z}_1 > \mathfrak{z}_2$, we retrieve on the most significant LWE the added value at the beginning of the algorithm. Otherwise, this added value is used during the subtraction. This value corresponds to $\text{Sign}(\mathfrak{z}_1 - \mathfrak{z}_2)$. By extracting this sign, and adding it to each LWE, with a PBS, we can see whether $\mathfrak{z}_1 > \mathfrak{z}_2$ or not and choose between the ciphertext obtain after the subtraction or its opposite and get $|\mathfrak{z}_1 - \mathfrak{z}_2|$.

Lemma 2 (Noise Constraints of Algorithm 2) The output noise variances of ciphertexts of Algorithm 2, \mathbf{ct}_s and $\mathbf{ct}_{\mathfrak{z}_{\text{sub}}}$, are respectively $4 \cdot \sigma_{\text{BR}}^2$ and σ_{BR}^2 .

To guarantee correctness of this operation, we need to find parameters that verify the following inequalities:

$$2 \cdot \sigma_{\text{in}}^2 + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2.$$

with σ_{BR} the noise added by the blind rotation, σ_{KS} the noise added by the keyswitch, σ_{CMUX} the noise added by the CMux and finally σ_{MS} , the noise added by the modulus switch.

Proof 2 (Proof of Lemma 2) Let us look at the noise propagation in Algorithm 2. We assume that each input ciphertext contains a noise following a centered Gaussian distribution with a variance σ_{in}^2 . These noises are also assumed to be statistically independent.

At the end of line 6, the variance of the noise in $\mathbf{ct}_{\mathfrak{z}_{\text{sub}}, i}$ is $2 \cdot \sigma_{\text{in}}^2$ for $0 \leq i < l_3$. The worst operation in terms of noise in the carry propagation

Algorithm 2: $(\mathbf{ct}_{\mathfrak{z}_{sub}}, \mathbf{ct}_{\text{sign}}) \leftarrow \text{IntSub}^* (\mathbf{ct}_{\mathfrak{z}_1}, \mathbf{ct}_{\mathfrak{z}_2}, \text{PUB})$

Context: $\begin{cases} \Delta_{\mathfrak{z}} : & \text{scaling factor} \\ \text{LUT}_{\text{Extract}} : & \text{LUT to extract the MSB (i.e. the } (\log_2(q) - 2)^{\text{th}} \text{ bit)} \\ \text{LUT}_{\text{f}} : & \text{LUT to return } \mathfrak{z} - q/4 \text{ if the MSB equals 1; } q/4 - \mathfrak{z} \text{ otherwise} \\ p_{\mathfrak{z}} : & \text{the carry-message modulus} \\ \beta_{\mathfrak{z}} : & \text{the message modulus} \end{cases}$

Input: $\begin{cases} \mathbf{ct}_{\mathfrak{z}_1} = [\mathbf{ct}_{\mathfrak{z}_1, \ell_{\mathfrak{z}}-1}, \dots, \mathbf{ct}_{\mathfrak{z}_1, 0}] \in \mathbb{Z}_q^{(n+1) \cdot \ell_{\mathfrak{z}}} \\ \mathbf{ct}_{\mathfrak{z}_2} = [\mathbf{ct}_{\mathfrak{z}_2, \ell_{\mathfrak{z}}-1}, \dots, \mathbf{ct}_{\mathfrak{z}_2, 0}] \in \mathbb{Z}_q^{(n+1) \cdot \ell_{\mathfrak{z}}} \\ \text{PUB} : \text{Public materials for } \mathbf{PBS\text{-}KS} \text{ and for } \mathbf{CBS} \end{cases}$

Output: $\begin{cases} \mathbf{ct}_{\mathfrak{z}_{sub}} = [\mathbf{ct}_{\mathfrak{z}_{sub}, \ell_{\mathfrak{z}}-1}, \dots, \mathbf{ct}_{\mathfrak{z}_{sub}, 0}] \in \mathbb{Z}_q^{(n+1) \cdot \ell_{\mathfrak{z}}} \\ \mathbf{ct}_{\mathfrak{s}} \in \mathbb{Z}_q^{n+1} \end{cases}$

```

1 for  $i$  in  $[1.. \ell_{\mathfrak{z}} - 1]$  do
    /* Ensure that  $\mathbf{ct}_{\mathfrak{z}_1, i}$  is larger than  $\mathbf{ct}_{\mathfrak{z}_2, i}$  */
2      $\mathbf{ct}_{\mathfrak{z}_1, i} \leftarrow \mathbf{ct}_{\mathfrak{z}_1, i} + \text{TrivialEncrypt}(2^{2 \cdot \rho_{\mathfrak{z}} - 1}, 1)$ 
3      $\mathbf{ct}_{\mathfrak{z}_1, i} \leftarrow \mathbf{ct}_{\mathfrak{z}_1, i} - \text{TrivialEncrypt}(2^{\rho_{\mathfrak{z}} - 1}, 1)$ 
4      $\mathbf{ct}_{\mathfrak{z}_{sub}, i} \leftarrow \mathbf{ct}_{\mathfrak{z}_1, i} - \mathbf{ct}_{\mathfrak{z}_2, i}$ 
5  $\mathbf{ct}_{\mathfrak{z}_1, 0} \leftarrow \mathbf{ct}_{\mathfrak{z}_1, 0} + \text{TrivialEncrypt}(2^{2 \cdot \rho_{\mathfrak{z}} - 1}, 1)$ 
6  $\mathbf{ct}_{\mathfrak{z}_{sub}, 0} \leftarrow \mathbf{ct}_{\mathfrak{z}_1, 0} - \mathbf{ct}_{\mathfrak{z}_2, 0}$ 
7  $\mathbf{ct}_{\mathfrak{z}_{sub}} \leftarrow \text{CarryPropagate}(\mathbf{ct}_{\mathfrak{z}_{sub}} = [\mathbf{ct}_{\mathfrak{z}_{sub}, \ell_{\mathfrak{z}}-1}, \dots, \mathbf{ct}_{\mathfrak{z}_{sub}, 0}], \text{PUB})$ 
    /* Extract the msb to get the sign */
8  $\mathbf{ct}_{\mathfrak{s}} \leftarrow \text{KS-PBS}(\mathbf{ct}_{\mathfrak{z}_{sub}, \ell_{\mathfrak{z}}-1}, \text{PUB}, \text{LUT}_{\text{Extract}})$ 
    /* Return the value if MSB==1, the opposite otherwise */
9 for  $i$  in  $[0.. \ell_{\mathfrak{z}} - 1]$  do
10      $\mathbf{ct}_{\mathfrak{z}_{sub}, i} \leftarrow \mathbf{ct}_{\mathfrak{z}_{sub}, i} + \mathbf{ct}_{\mathfrak{s}}$ 
11      $\mathbf{ct}_{\mathfrak{z}_{sub}, i} \leftarrow \text{KS-PBS}(\mathbf{ct}_{\mathfrak{z}_{sub}, i}, \text{PUB}, \text{LUT}_{\text{f}})$ 
12  $\mathbf{ct}_{\mathfrak{z}_{sub}, \ell_{\mathfrak{z}}-1} \leftarrow \text{KS-PBS}(\mathbf{ct}_{\mathfrak{z}_{sub}, \ell_{\mathfrak{z}}-1}, \text{PUB}, \text{LUT}_{\text{f}})$ 
    /* Put the sign on the padding bit plus flip the bit to keep
       the representation 0 is positive and 1 is negative */
13  $\mathbf{ct}_{\mathfrak{s}} \leftarrow \mathbf{ct}_{\mathfrak{s}} \cdot 2 + \text{TrivialEncrypt}(q/2, 1)$ 
14 return  $(\mathbf{ct}_{\mathfrak{s}}, \mathbf{ct}_{\mathfrak{z}_{sub}} = [\mathbf{ct}_{\mathfrak{z}_{sub}, \ell_{\mathfrak{z}}-1}, \dots, \mathbf{ct}_{\mathfrak{z}_{sub}, 0}])$ 

```

on line 7 consists of adding a freshly bootstrapped ciphertext with one of the $\mathbf{ct}_{\mathfrak{z}_{sub}, i}$ and applying to it a keyswitch and a bootstrapping. It means that to

have correctness we must verify the following inequality:

$$2 \cdot \sigma_{\text{in}}^2 + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2.$$

with $\sigma_{\text{BR}}^2, \sigma_{\text{KS}}^2, \sigma_{\text{MS}}^2$, the noise after respectively a bootstrapping, a keyswitch and a modulus switch and t^2 , the noise bound as defined in the proof of the noise constraints of Algorithm 6.

On line 8, another noise constraint appears.

$$\sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2.$$

As the left hand term is smaller than the one of the previous inequality, we can discard this constraint. On lines 11 and 12, we have the following constraints

$$2 \cdot \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2 \ \& \ \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2.$$

Using the fact that the last constraint is dominated by the others, we can remove it from the set of constraints. In the end of Algorithm 2, the sign ct_s has a noise variance of $4 \cdot \sigma_{\text{BR}}^2$ and the each ciphertext in the vector ct_{sub} has a noise variance σ_{BR}^2 .

2.5 Traditional Floating-Point Representation

Floating-points have become the standard to represent real numbers in computer science, as described in [Ins08]. Their main advantage lies in having a variable precision all along computations, giving more flexibility and accuracy. Usually a floating-point is represented by three values: the sign, the mantissa and the exponent. The most common floating-point encodings on CPU are the single precision, represented on 32 bits (with 1 bit of sign, 8 bits of exponent and 23 bits of mantissa) and the double precision over 64 bits (1 bit of sign, 11 bits of exponent and 52 bits of mantissa). Less common but still useful encodings are the half-precision which represents 16-bit floats, or some alternative called Bfloat [WK19]. Without getting into details, these encodings mainly differ in the distribution of the bit number associated to the mantissa and the exponent. Finally, another family called MiniFloat is dedicated to floats whose the size is 8 bits. We refer to [MBDD⁺18] for more information about floats.

Definition 4 (Floating-Point) Let $\mathfrak{b} \in \mathbb{N}$ such that $\mathfrak{b} \geq 2$. Let $\max_{\epsilon} \in \mathbb{N}^*$ and a fixed $\text{bias} \in [0, \max_{\epsilon}]$. A floating-point number $x \in \mathbb{R}$ is partially characterized by three values $(\mathfrak{s}, \mathfrak{m}, \mathfrak{e}) \in \{0; 1\} \times \mathbb{N} \times [0, \max_{\epsilon}]$, such that: $x = (-1)^{\mathfrak{s}} \cdot \mathfrak{m} \cdot \mathfrak{b}^{\mathfrak{e} - \text{bias}}$. With this definition, a floating-point number may have several representations. To obtain a unique representation, the mantissa \mathfrak{m} must be in the interval $[1, \mathfrak{b})$ except for the value zero where $\mathfrak{m} = 0$.

3 Homomorphic Floating-Points (HFP)

In this section, we present a promising but limited approach in terms of precision which leverages the WoP-PBS to efficiently compute any operation over floating-point numbers. Then, we explain how to efficiently translate the traditional representation of floating points into a TFHE-friendly way for larger precisions. Finally, we describe the first building blocks needed to perform higher-level operations on these homomorphic floating-points.

3.1 MiniFloats: WoP-PBS Based Floats

A powerful approach to defining homomorphic floats for TFHE-like schemes relies on the WoP-PBS. The method is somewhat similar to the gate bootstrapping approach defined in [CGGI20]: almost every operation is performed using a WoP-PBS.

Minifloat Encoding Let ρ be the number of bits of precision for a message in an LWE ciphertext, and let $\rho_{\mathfrak{m}}$ (resp. ρ_{ϵ}) be the number of bits of the mantissa (resp. the exponent). In this first attempt at building TFHE-minifloats, we do not need to have distinct ciphertexts for the mantissa, the sign and the exponent. For instance, we can define an 8-bit floating-point with $\rho = 4$, $\rho_{\mathfrak{m}} = 3$ and $\rho_{\epsilon} = 4$ using only two LWEs, $\text{ct}_1 \in \text{LWE}_s(\mathfrak{s} || \mathfrak{m}_2 || \mathfrak{m}_1 || \mathfrak{m}_0)$ and $\text{ct}_2 \in \text{LWE}_s(\mathfrak{e}_3 || \mathfrak{e}_2 || \mathfrak{e}_1 || \mathfrak{e}_0)$ where \mathfrak{m}_i (resp. \mathfrak{e}_i) corresponds to the i^{th} bit of the mantissa (resp. of the exponent). Each element can be dispatched in any order, but the order must be publicly known to correctly generate the LUT. We call this encoding the *minifloat encoding* and we write it as follows:

$$\text{TFHE-Minifloat}^{\rho}(\rho_{\mathfrak{m}}, \rho_{\epsilon}, \text{bias}) = \text{Encoding}^{\rho}(\mathfrak{s} || \mathfrak{m} || \mathfrak{e} ||).$$

Minifloat Operations Defining operations over the minifloat encoding is easy: each one of them is computed with a WoP-PBS where the LUT associated with the operation is given. The WoP-PBS can easily be extended to take many ciphertexts as input in order to compute multivariate operations: the bit extraction step can be done for every input and then a single CMux tree using the bits of every input. Let Op be an operation (e.g., an addition), and LUT_{Op} its associated LUT (for more details on how to properly generate the LUT, we refer to [BBB⁺22]). For some $k \in \mathbb{N}$, let ct_{f_i} for $i \in [0, k - 1]$ be the input ciphertexts. Then, the output is given by: $\text{ct}_{f_{\text{out}}} \leftarrow \text{WoP-PBS}(\{\text{ct}_{f_i}\}_{i=0}^{k-1}, \text{LUT}_{\text{Op}}, \text{PUB})$.

The main advantage is about the complexity of this method, which is not dependent on the functions that have to be computed, i.e., any univariate functions will take the same time (e.g., cosine, logarithm, ...).

Remark 4 *Some operations do not require a complete WoP-PBS. For example, to perform a ReLU, we only need to extract the sign and perform a CMux between input value and a trivial LWE (defined in 4) that encrypts zero.*

We provide benchmarks for this method in Section 6. This method is very efficient but it is limited in terms of precision. In fact, this method does not work when the combined bit size of all inputs of a WoP-PBS exceeds approximately twenty bits, because the number of values that need to be represented is too large and the LUT quickly becomes too big [BBB⁺22, Remark 8]. Next section explores another encoding that can efficiently support large floating-point numbers.

3.2 Homomorphic Floating-Point Encoding

As in the traditional representation of floating-point numbers, the homomorphic floating point representation is divided into three parts: the sign, the mantissa and the exponent.

The **sign** (\mathfrak{s}) is encoded by one LWE ciphertext. This ciphertext encrypts the value 0 if the sign is positive or 1 if the sign is negative.

The **mantissa** (\mathfrak{m}) is encoded by several LWE ciphertexts (at least 2). Each ciphertext associated with the mantissa encodes the same amount of message bits (denoted $\rho_{\mathfrak{m}}$ in the following). For a mantissa represented by $\ell_{\mathfrak{m}}$ LWE ciphertexts, we can represent integers in $[0; 2^{\ell_{\mathfrak{m}} \cdot \rho_{\mathfrak{m}}})$. The ciphertext encrypting the most significant bits (respectively, the least significant bits) of

the mantissa is called the most significant (respectively, the least significant) ciphertext. With this representation, we can ensure that the precision of the mantissa is at least $((\ell_m - 1) \cdot \rho_m + 1)$ -bits. Indeed, the least significant ciphertext will be discarded after some operations, so the information in this block must not be seen as additional precision: when the carry space of the most significant **LWE** ciphertext is full, a new **LWE** ciphertext is added as the new most significant block. The less significant block is then removed and the exponent value is increased. In floating point arithmetic, this approach is generally called *rounding towards zero*. This means that the least significant bits from the mantissa are discarded. This way, the exponent can be smaller and represent a large range of values. In our representation, to keep a unique encoding for any floating-point, the most significant block should always be different from zero (except for the special value zero where all the blocks are equal to zero). So for any non-zero values, the mantissa is an integer in $[2^{\rho_m \cdot (\ell_m - 1)}; 2^{\rho_m \cdot \ell_m})$.

The **exponent** (ϵ) is encoded by one or more **LWE** ciphertexts. Each **LWE** ciphertext encrypts the same amount of bits ρ_ϵ . The value represented by the exponent is in base 2^{ρ_m} (as already mentioned in the mantissa). So an exponent encrypted in ℓ_ϵ **LWE** ciphertexts represents values in $[0; (2^{\rho_m})^{2^{\ell_\epsilon \cdot \rho_\epsilon}})$. To represent an exponent that can be negative or positive, the positive value encoded in these ℓ_ϵ **LWE** ciphertexts needs to be subtracted by a value named **bias**. When we decode, we obtain, $e \in [(2^{\rho_m})^{-\text{bias}}; (2^{\rho_m})^{2^{\ell_\epsilon \cdot \rho_\epsilon} - \text{bias}})$. The **encoding of the TFHE floating-point** is illustrated in Figure 1, and we refer to it as follows:

$$\text{TFHE_FP}(\ell_m, \rho_m, \ell_\epsilon, \rho_\epsilon, \text{bias}).$$

Bias The value **bias** can be set to any value, but to represent a large range of values, in the traditional floating-point, this value is often set to be half of the range of the exponent. With our representation, this value should correspond to $2^{\ell_\epsilon \cdot \rho_\epsilon - 1}$, but in our algorithm we choose to use $\text{bias} = 2^{\ell_\epsilon \cdot \rho_\epsilon - 1} + \ell_m - 1$. Through this specific value, we gain a speed-up in the homomorphic floating multiplication proposed in Algorithm 10.

Special Values In the floating-point arithmetic, the subnormal values are the closest values to zero: they are represented by an exponent equal to zero and the leading significant digits equal to zero. In our implementation

we choose to not represent this values to have better performance, but our algorithm can be easily modified to take these values into account. In what follows, the leading significant block is always strictly positive except for the zero value. This is the only value represented by each mantissa and exponent blocks equals to zero. Thus, if an operation yields an encrypted float which has its most significant **LWE** equal to zero, the result will encrypt the zero value (i.e., all the mantissa and exponent **LWE** will be equal to zero). To keep the algorithms easier to read, other special values like infinities, or NaN are voluntary excluded. Note that the process to support these is nevertheless detailed in Section 7.1.

Encoding and Encrypting We propose an algorithm to encode and decode real numbers for TFHE with the representation $\text{TFHE_Fp}(\ell_m, \rho_m, \ell_e, \rho_e, \text{bias})$. Let f be a real number. First, we need to find $m \in [0, \ell_m \cdot \rho_m)$ and $e \in [0, \ell_e \cdot \rho_e)$ such that:

$$f = (-1)^s \cdot m \cdot (2^{\rho_m})^{e - \text{bias}}.$$

To obtain a unique representation of these floating-point representations, we impose that the most significant block of the mantissa must be strictly positive (except for the value zero). From this first encoding, we split the mantissa and the exponent according to the 2^{ρ_m} and 2^{ρ_e} -radix decompositions, i.e., $m \leftarrow \text{Encoding}^{2^{\rho_m}}(m)$ and $e \leftarrow \text{Encoding}^{2^{\rho_e}}(e)$. The final encoding is given by:

$$f = (s, m = (m_{\ell_m-1}, \dots, m_0)_{\rho_m}, e = (e_{\ell_e-1}, \dots, e_0)_{\rho_e}).$$

In absolute value, the maximum value represented by this encoding is $\max = (2^{\ell_m \cdot \rho_m} - 1) \cdot (2^{\rho_m})^{2^{\ell_e \cdot \rho_e} - \text{bias} - 1}$. Since the subnormal values are not taken into account, the minimum positive value reached by this encoding (without zero) is $\min = (2^{\ell_m \cdot \rho_m - 1}) \cdot (2^{\rho_m})^{-\text{bias}}$. In Algorithm 3 (resp. Algorithm 4), the method to encrypt (resp. decrypt) a floating-point number is described. The following lemma states the correctness and the notations used to represent homomorphic floats.




Figure 1: The figure illustrates the encoding of a homomorphic floating-point. The Mantissa (in pink) and the Exponent (in green) are split in several ciphertexts which each ciphertext encrypting 4 bits. The fully colored boxes in the figure represent the bits of messages space and the empty colored boxes correspond to the bits of carry space. The sign (in blue) is encoded in only one ciphertext where the information is encrypted on the most significant bit.

Algorithm 3: $\text{ct}_f \leftarrow$	Algorithm 4: $f \leftarrow$
$\text{EncryptFloat}(s, m, e)$	$\text{DecryptFloat}(\text{ct}_f)$
Input: $\text{EncodeFloat}(x) = (s, m, e)$ Output: $\text{ct}_f = [\text{ct}_s, \text{ct}_m, \text{ct}_e]$ 1 $\text{ct}_s \leftarrow \text{Encrypt}_s(\frac{s}{2} \cdot s) \in \text{LWE}(s)$ 2 $\text{ct}_m = (\text{ct}_{m, \ell_m-1}, \dots, \text{ct}_{m,0}) \leftarrow \text{Encrypt}_s(m)$ 3 $\text{ct}_e = (\text{ct}_{e, \ell_e-1}, \dots, \text{ct}_{e,0}) \leftarrow \text{Encrypt}_s(e)$ 4 return $\text{ct}_f = [\text{ct}_s, \text{ct}_m, \text{ct}_e]$	Input: $\text{ct}_f = [\text{ct}_s, \text{ct}_m, \text{ct}_e]$ Output: $f = (-1)^s \cdot m \cdot (2^{\rho_m})^{e-\text{bias}} \in \mathbb{R}$ 1 If $\text{Decrypt}_s(\text{ct}_s) = 0$ then $s = 1$, Else $s = -1$ 2 $m \leftarrow \text{Decrypt}_s(\text{ct}_m)$ 3 $e \leftarrow \text{Decrypt}_s(\text{ct}_e)$ 4 return $f \leftarrow (-1)^s \cdot m \cdot (2^{\rho_m})^{e-\text{bias}}$

Lemma 3 (Correctness of DecryptFloat (4)) *Let $f = (s, m, e) \in \{0, 1\} \times [0, \ell_m \cdot \rho_m] \times [0, \ell_e \cdot \rho_e]$. Let $\text{ct}_f \leftarrow \text{EncryptFloat}(f)$ such that $\text{ct}_f = [\text{ct}_s, \text{ct}_e, \text{ct}_m]$, with $\text{ct}_s \in \text{LWE}_s(s)$, $\text{ct}_e = [\text{ct}_{e, \ell_e-1}, \dots, \text{ct}_{e,0}] \in [\text{LWE}_s(e_{1_{\ell_e-1}}), \dots, \text{LWE}_s(e_{1_0})]$ and $\text{ct}_m = [\text{ct}_{m, \ell_m-1}, \dots, \text{ct}_{m,0}] \in [\text{LWE}_s(m_{1_{\ell_m-1}}), \dots, \text{LWE}_s(m_{1_0})]$. Then, $\text{DecryptFloat}_s(\text{ct}_f) = f$.*

Trivial Encrypt A trivial encryption is an LWE ciphertext where all the a_i are equal zero (for $i \in [0, n]$). This is trivially extendable to the floats. This is denoted $\text{TrivialEncryptFloat}(f)$. Sometimes, we only need to trivially encrypt a part of a floating-point: which is suggested by the notation $\text{TrivialEncrypt}(\text{Value}, \text{NumberOfBlocks})$. For instance, to encrypt the value v as an exponent, we note $\text{TrivialEncrypt}(v, \ell_e)$.

Definition 5 (Maximum error after operation) *In the context of floating-point, due to the encoding, after each operation a small error could be introduced (this error is not tied to the TFHE noise). This added error is denoted ϵ . To find the errors added after each operation by our encoding, we look at the maximal error added in the mantissa and we multiply this error by the exponent. For a floating-point $\mathfrak{f} = (-1)^s \cdot \mathfrak{m} \cdot (2^{\rho_m})^{\epsilon - \text{bias}}$, the error after an operation can be bounded by $\text{error}_m \cdot (2^{\rho_m})^{\epsilon - \text{bias}}$. As we do not represent the subnormal values, if the ϵ is equal to 0, the error is bounded by $2^{\rho_m \cdot (\ell_m - 1)} \cdot (2^{\rho_m})^{-\text{bias}}$.*

Example: Encoding a 64 bits Floating-Point In the [Ins08] standard, a floating-point is composed of 1 bit of sign, 11 bits of exponent and 52 bits plus one hidden bit of mantissa. So as mentioned in the beginning of Section 3, to ensure a precision of at least 53 bits, we need to have $(\ell_m - 1) \cdot \rho_m + 1 \geq 53$ i.e., one additional block to perform operations without losing the precision of 53 bits. For the mantissa, we choose $\ell_m = 27$ with $\rho_m = 2$. In [Ins08], the exponent value ϵ belongs to $[-1023, 1024)$. To simplify the implementation, we prefer to have $\rho_\epsilon = \rho_m$, and a bias equal to $2^{\ell_\epsilon \cdot \rho_\epsilon - 1} + \ell_m - 1$ with $\ell_\epsilon = 5$. Thus, this yields in a floating-point exponent in $[-\text{bias} \cdot \rho_m, (2^{\ell_\epsilon \cdot \rho_\epsilon} - \text{bias} - 1) \cdot \rho_m] = [-1076, 970]$ with $\text{bias} = 538$. This allows representing as many values as the standard one, but with a different scale: the upper bound is lower, but more precision is given near values close to zero. These parameters correspond to the representation :

$$\text{TFHE_Fp}_{64b}(\ell_m = 27, \rho_m = 2, \ell_\epsilon = 5, \rho_\epsilon = 2, \text{bias} = 538).$$

More encoding examples for 32-bits, 16-bits and 8-bits are given in Table 3 (Section 6).

3.3 Choosing Between Two Ciphertexts

In what follows, we extensively use Algorithm 5 to homomorphically make a choice between two LWE ciphertext lists depending on an encrypted bit. This algorithm is an extended version of the CMux described in [CGGI20, Lemma 3.16]. The selector is a GGSW ciphertext, and the choice is done between two lists of LWE ciphertexts.

Lemma 4 *Let $\mathbf{ct}_i = [\mathbf{ct}_{\mathfrak{z}_i, \ell_3 - 1}, \dots, \mathbf{ct}_{\mathfrak{z}_i, 0}] \in [\text{LWE}_s(\mathfrak{z}_i, \ell_3 - 1), \dots, \text{LWE}_s(\mathfrak{z}_i, 0)] \subseteq \mathbb{Z}_q^{(n+1) \cdot \ell_3}$ with $i \in \{0, 1\}$ be two ciphertexts encrypting $\mathfrak{z}_i \in \mathbb{N}$. Let $\overline{\text{CT}}_{\text{Sel}} \in$*

Algorithm 5: $\mathbf{ct}_{\text{res}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_0, \mathbf{ct}_1, \overline{\overline{\mathbf{CT}}}_{\text{Sel}})$

Input: $\begin{cases} \mathbf{ct}_0 = [\text{ct}_{0,\ell_3-1}, \dots, \text{ct}_{0,0}] \in \mathbb{Z}_q^{(n+1) \cdot \ell_3} \text{ with } \text{ct}_{i,j} \in \text{LWE}_s(m_{i,j}) \text{ for some plaintext } m_{i,j} \\ \mathbf{ct}_1 = [\text{ct}_{1,\ell_3-1}, \dots, \text{ct}_{1,0}] \in \mathbb{Z}_q^{(n+1) \cdot \ell_3} \\ \overline{\overline{\mathbf{CT}}}_{\text{Sel}} \in \text{GGSW}_S^{\beta,\ell}(b) \text{ (with } b \in \{0, 1\}) \end{cases}$

Output: $\begin{cases} \mathbf{ct}_{\text{res}} = [\text{ct}_{\text{res},\ell_3-1}, \dots, \text{ct}_{\text{res},0}] \in \mathbb{Z}_q^{(n+1) \cdot \ell_3} \end{cases}$

```

1 for  $i$  in  $[0..\ell_3)$  do
2    $\text{CT}_{0,i} \leftarrow \text{SampleInsert}(\text{ct}_{0,i})$ ,  $\text{CT}_{1,i} \leftarrow \text{SampleInsert}(\text{ct}_{1,i})$ 
3    $\text{CT}_{\text{res},i} \leftarrow \text{CMux}(\text{CT}_{0,i}, \text{CT}_{1,i}, \overline{\overline{\mathbf{CT}}}_{\text{Sel}})$ 
4    $\text{ct}_{\text{res},i} \leftarrow \text{Sample\_extract}(\text{CT}_{\text{res},i})$ 
5 return  $\mathbf{ct}_{\text{res}} = [\text{ct}_{\text{res},\ell_3-1}, \dots, \text{ct}_{\text{res},0}]$ 

```

$\text{GGSW}_S^{\beta,\ell}(b)$ (with $b \in \{0, 1\}$). Let $\mathbf{ct}_{\text{res}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_0, \mathbf{ct}_1, \overline{\overline{\mathbf{CT}}}_{\text{Sel}})$. Then, $\text{Decrypt}_s(\mathbf{ct}_{\text{res}}) = \mathfrak{z}b$.

Proof 3 (Correctness of Algorithm 5) Let us execute each instruction of Algorithm 5. In Lines 2 and 3, both LWE ciphertexts are transformed into GLWE ciphertexts with the message on the first coefficient and random messages on all the other coefficients, i.e., $\text{CT}(M_{j,i}) \in \text{GLWE}_s(M_{j,i})$ with $M_{j,i} = m_{j,i} + \sum_{\alpha=1}^{N-1} r_{j,i,\alpha} X^\alpha$, for some $r_{j,i,\alpha} \in \mathbb{Z}_q$ with $j \in \{0, 1\}$. Next, in Line 4: $\text{CT}_{\text{res},i}$ is the result of the CMux, i.e., $\text{CT}_{\text{res},i}(M_{b,i}) \leftarrow \text{CMux}(\text{CT}(M_{1,i}), \text{CT}(M_{0,i}), \overline{\overline{\mathbf{CT}}}_{\text{Sel}}(b))$, (as define in 2.3), with $b \in \{0, 1\}$. At Line 5, we extract the first coefficient of $\text{CT}_{\text{res},i}$. The result is then $\text{ct}(m_{\text{res},i}) \in \text{LWE}_s(m_{b,i})$.

Proof 4 (CMux Noise Analysis) We adapt [CLOT21, Proof 11] for an external product using a binary secret in the GGSW ciphertext. The only difference is that the GGSW ciphertext is obtained through a CBS rather than a freshly encrypted ciphertext, so we use the same noise formula. However, the noise of the bootstrapping key is defined as the output noise of circuit bootstrapping instead of fresh encryption.

3.4 Propagating the Carries

Since HFP are based on radix-based homomorphic integers, the need to propagate the carry must be considered to ensure correctness. Indeed, in each

block, carry might accumulate all along computation up to point where the carry space is full. Differently from modular integer computation, where the modulus is generally a power of two, we cannot simply remove the carry from the most significant block. In our case, when the mantissa has a carry that has been propagated to the most significant block, a new one must be created and the last one can be removed. In Algorithm 6, we describe the process to perform this homomorphically. It takes as input a ciphertext encrypting a floating-point, and returns another ciphertext where the carries have been propagated.

Algorithm 6: $\mathbf{ct}_{\text{f}_{\text{out}}} \leftarrow \text{CarryPropagateFloat}(\mathbf{ct}_{\text{f}})$

Context: LUT_{id} : Lookup Table associated to the id function.

Input: $\begin{cases} \mathbf{ct}_{\text{s}} \in \text{LWE}_{\text{s}}(\text{s}) \\ \mathbf{ct}_{\text{e}} = [\text{ct}_{\text{e}, \ell_{\text{e}}-1}, \dots, \text{ct}_{\text{e}, 0}] \in [\text{LWE}_{\text{s}}(\text{e}_{\ell_{\text{e}}-1}), \dots, \text{LWE}_{\text{s}}(\text{e}_0)] \\ \mathbf{ct}_{\text{m}} = [\text{ct}_{\text{m}, \ell_{\text{m}}-1}, \dots, \text{ct}_{\text{m}, 0}] \in [\text{LWE}_{\text{s}}(\text{m}_{\ell_{\text{m}}-1}), \dots, \text{LWE}_{\text{s}}(\text{m}_0)] \\ \text{PUB} : \text{Public keys for KS-PBS and for CBS} \end{cases}$

Output: $\begin{cases} \mathbf{ct}_{\text{s}_{\text{out}}} \in \text{LWE}_{\text{s}}(\text{s}) \\ \mathbf{ct}_{\text{e}_{\text{out}}} = [\text{ct}_{\text{e}_{\text{out}}, \ell_{\text{e}}-1}, \dots, \text{ct}_{\text{e}_{\text{out}}, 0}] \in [\text{LWE}_{\text{s}}(\text{e}_{\text{out}, \ell_{\text{e}}-1}), \dots, \text{LWE}_{\text{s}}(\text{e}_{\text{out}, 0})] \\ \mathbf{ct}_{\text{m}_{\text{out}}} = [\text{ct}_{\text{m}_{\text{out}}, \ell_{\text{m}}-1}, \dots, \text{ct}_{\text{m}_{\text{out}}, 0}] \in [\text{LWE}_{\text{s}}(\text{m}_{\text{out}, \ell_{\text{m}}-1}), \dots, \text{LWE}_{\text{s}}(\text{m}_{\text{out}, 0})] \end{cases}$

```

1  $\mathbf{ct}_{\text{s}_{\text{out}}} \leftarrow \text{KS-PBS}(\mathbf{ct}_{\text{s}}, \text{LUT}_{\text{id}}, \text{PUB})$ 
2  $\mathbf{ct}_{\text{e}'} \leftarrow \text{CarryPropagate}(\mathbf{ct}_{\text{e}}, \text{PUB})$ 
3  $\mathbf{ct}_{\text{m}'} = [\text{ct}_{\text{m}', \ell_{\text{m}}}, \dots, \text{ct}_{\text{m}', 0}] \leftarrow \text{CarryPropagateExtend}(\mathbf{ct}_{\text{m}}, \text{PUB});$ 
   /* Algo 1 */
4  $\mathbf{ct}_{\text{m}_{+}} = [\text{ct}_{\text{m}', \ell_{\text{m}}}, \dots, \text{ct}_{\text{m}', 1}], \mathbf{ct}_{\text{m}_{-}} = [\text{ct}_{\text{m}', \ell_{\text{m}}-1}, \dots, \text{ct}_{\text{m}', 0}]$ 
5  $\overline{\text{CT}} \leftarrow \text{CBS}(\text{ct}_{\text{m}'_{\ell_{\text{m}}}}, \text{PUB})$ 
6  $\mathbf{ct}_{\text{m}_{\text{out}}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{\text{m}_{-}}, \mathbf{ct}_{\text{m}_{+}}, \overline{\text{CT}}),$ 
    $\mathbf{ct}_{\text{e}_{\text{out}}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{\text{e}'}, \mathbf{ct}_{\text{e}'} + \text{TrivialEncrypt}(1, \ell_{\text{e}}), \overline{\text{CT}});$  /* Algo 5
   */
7 return  $\mathbf{ct}_{\text{f}_{\text{out}}} = (\mathbf{ct}_{\text{s}_{\text{out}}}; \mathbf{ct}_{\text{e}_{\text{out}}}; \mathbf{ct}_{\text{m}_{\text{out}}})$ 

```

Lemma 5 Let $\mathbf{ct}_{\text{f}} = [\mathbf{ct}_{\text{s}}, \mathbf{ct}_{\text{m}}, \mathbf{ct}_{\text{e}}]$ with $\mathbf{ct}_{\text{s}} \in \text{LWE}_{\text{s}}(\text{s})$, $\mathbf{ct}_{\text{e}} = [\text{ct}_{\text{e}, \ell_{\text{e}}-1}, \dots, \text{ct}_{\text{e}, 0}] \in [\text{LWE}_{\text{s}}(\text{e}_{\ell_{\text{e}}-1}), \dots, \text{LWE}_{\text{s}}(\text{e}_0)]$, $\mathbf{ct}_{\text{m}} = [\text{ct}_{\text{m}, \ell_{\text{m}}-1}, \dots, \text{ct}_{\text{m}, 0}] \in [\text{LWE}_{\text{s}}(\text{m}_{\ell_{\text{m}}-1}), \dots, \text{LWE}_{\text{s}}(\text{m}_0)]$ such that some $\text{ct}_{\text{m}, i}$ (resp. $\text{ct}_{\text{e}, i}$) encrypting some $\text{m}_i > 2^{\rho_{\text{m}}}$ (resp. $\text{e}_i > 2^{\rho_{\text{e}}}$). Let $\mathbf{ct}_{\text{f}_{\text{out}}} \leftarrow \text{CarryPropagateFloat}(\mathbf{ct}_{\text{f}})$. Let $\text{f} = \text{DecryptFloat}_{\text{s}}(\mathbf{ct}_{\text{f}}) =$

$(-1)^s \cdot m \cdot 2^{\rho_m \epsilon - \text{bias}}$ and $f_{\text{out}} = \text{DecryptFloat}_s(\mathbf{ct}_{f_{\text{out}}}) = (-1)^{s_{\text{out}}} \cdot m_{\text{out}} \cdot 2^{\rho_m \epsilon_{\text{out}} - \text{bias}}$.
 If $m \in [2^{\rho_m \cdot (\ell_m - 1)}, 2^{\rho_m \cdot \ell_m}]$, then $m_{\text{out}} \in [2^{\rho_m \cdot (\ell_m - 1)}, 2^{\rho_m \cdot \ell_m}]$. Else if $m \in [2^{\rho_m \cdot \ell_m}, 2^{\rho_m \cdot (\ell_m + 1)}]$, then $m_{\text{out}} \in [2^{\rho_m \cdot (\ell_m - 1)}, 2^{\rho_m \cdot \ell_m}]$ and $\epsilon_{\text{out}} = \epsilon + 1$.
 Moreover, for all $i \in [0, \ell_m - 1]$ (resp. $i \in [0, \ell_\epsilon - 1]$), $m_{\text{out},i} \in [0, 2^{\rho_m} - 1]$ (resp. $\epsilon_{\text{out},i} \in [0, 2^{\rho_\epsilon} - 1]$).

Proof 5 (Correctness of CarryPropagateFloat (Algorithm 6)) In Line 2, we get $\mathbf{ct}_{\epsilon'} = [\mathbf{ct}_{\epsilon', \ell_\epsilon - 1}, \dots, \mathbf{ct}_{\epsilon', 0}]$ such that $\forall i \in [0, \ell_\epsilon - 1]$, $\text{Decrypt}(\mathbf{ct}_{\epsilon', i}) = \epsilon'_i < 2^{\rho_\epsilon}$ after the carry propagation. Likewise, in Line 3, we do an extended carry propagation of the mantissa, so that $\mathbf{ct}_{m'} = [\mathbf{ct}_{m', \ell_m}, \dots, \mathbf{ct}_{m', 0}]$ such that, $\forall i \in [0, \ell_m]$, $\text{Decrypt}(\mathbf{ct}_{m', i}) = m'_i < 2^{\rho_m}$. Note that the carry on the most significant block is not lost and create a new LWE ciphertext \mathbf{ct}_{m', ℓ_m} encrypting the propagated carry of $\mathbf{ct}_{m, \ell_m - 1}$.

In the next steps, the idea is to decide if we need to keep this block \mathbf{ct}_{m', ℓ_m} and remove the least significant block (i.e., return \mathbf{ct}_{m_+}) or if we can discard it (i.e., return \mathbf{ct}_{m_-}). This allows us to output a result which has the same number of blocks ℓ_m than the input. To do so, in Line 6 we perform a circuit bootstrapping returning a GGSW ciphertext: $\overline{\overline{\mathbf{CT}}} \in \text{GGSW}_S^{\beta, \ell}(0)$ if the new \mathbf{ct}_{m', ℓ_m} is in $\text{LWE}_s(0)$, otherwise, $\overline{\overline{\mathbf{CT}}}$ is in $\text{GGSW}_S^{\beta, \ell}(1)$. Next in Line 7, Algorithm 5 returns \mathbf{ct}_{m_+} if $\overline{\overline{\mathbf{CT}}}$ is in $\text{GGSW}_S^{\beta, \ell}(1)$, or \mathbf{ct}_{m_-} otherwise. Likewise, in the case where \mathbf{ct}_{m', ℓ_m} does not encrypt 0, the exponent should be updated. The condition is then the same as previously, so that we can use the same selector $\overline{\overline{\mathbf{CT}}}$ to choose between the initial value of the exponent $\mathbf{ct}_{\epsilon'}$ or the one which has been increased by one $\mathbf{ct}_{\epsilon'} + \text{TrivialEncrypt}(1, \ell_\epsilon)$.

Remark 5 (Carry Propagation & Refresh) After most operations, we will apply Algorithm 6 to properly propagate the carries and refresh the noise. However, after operations like the ReLU (Algorithm 12) or the approximated Sigmoid (Algorithm 13) that do not fill the carry block, we only need to perform a PBS on each ciphertext to obtain fresh noise.

Lemma 6 (Noise Constraints of Algorithm 6) The output ciphertexts of Algorithm 6, $\mathbf{ct}_{m_{\text{out}}}$ has a noise variance $\sigma_{\text{BR}}^2 + \sigma_{\text{CMux}}^2$, $\mathbf{ct}_{\epsilon_{\text{out}}}$ has a noise variance $\sigma_{\text{BR}}^2 + \sigma_{\text{CMux}}^2$ and $\mathbf{ct}_{s_{\text{out}}}$ has a noise variance σ_{BR}^2 .

To guarantee correctness of Algorithm 6, we need to find parameters that verify the following inequalities:

$$\sigma_{\text{in}, \epsilon}^2 + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2 \quad \& \quad \sigma_{\text{in}, m}^2 + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2.$$

With $\sigma_{\text{in},\epsilon}$ the noise variance of the input exponent ciphertexts, $\sigma_{\text{in},m}$ the noise variance of the input mantissa ciphertexts and with σ_{BR} the noise added by the blind rotation, σ_{KS} the noise added by the keyswitch, σ_{CMux} the noise added by the CMux and finally σ_{MS} , the noise added by the modulus switch.

Proof 6 (Proof of Lemma 6) In this proof, we use the same techniques as those introduced in [BBB⁺22]. In particular, we use the noise bound [BBB⁺22, Def. 8], a quantity representing the maximum noise variance that still guarantees the correctness up to some failure probability. We also use a simpler version of [BBB⁺22, Theorem 1] which consists of removing redundant inequalities and dominated constraints. Simply put, if we have two inequalities $f(x) + g(x) \leq t$ and $f(x) \leq t$ with f and g two positive functions, we can focus on the first one as the second one will be automatically satisfied if the first is. In this proof and the next ones, when this situation arises, we will say that the second inequality is dominated by the first.

Let us assume that the input ciphertexts ct_s , ct_ϵ and ct_m have respectively the following noise variances $\sigma_{\text{in},s}^2$, $\sigma_{\text{in},\epsilon}^2$ and $\sigma_{\text{in},m}^2$. The first line of the algorithm consists in a keyswitch and a bootstrapping. We have the following noise constraint: $\sigma_{\text{in},s}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$, with σ_{KS}^2 and σ_{MS}^2 the noise added respectively by the keyswitch and the modulus switch. t is a noise bound such that $t = \frac{\Delta}{2 \cdot z^*(p_{\text{fail}})}$ with the standard score $z^*(p_{\text{fail}}) = \sqrt{2} \cdot \text{erf}^{-1}(1 - p_{\text{fail}})$ and the scaling factor Δ introduced in Definition 1. If we find parameters that guarantee the inequality above, the bootstrapping will be successful with probability $1 - p_{\text{fail}}$.

Then, we have a carry propagate and an extended carry propagate. We refer to the analysis for Algorithm 2 for the explanation about the constraints in these algorithms:

$$\sigma_{\text{in},\epsilon}^2 + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2 \quad \& \quad \sigma_{\text{in},m}^2 + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2.$$

Finally, we have a circuit bootstrapping that also creates a noise constraint $\sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$. Notice that the left-hand side here is smaller than in both inequalities above, so we can remove this last inequality from the set of constraints. Then, the output ciphertexts $\text{ct}_{m,\text{out}}$, $\text{ct}_{\epsilon,\text{out}}$ and $\text{ct}_{s,\text{out}}$ have respectively a noise variance $\sigma_{\text{BR}}^2 + \sigma_{\text{cmux}}^2$, $\sigma_{\text{BR}}^2 + \sigma_{\text{cmux}}^2$ and σ_{BR}^2 with σ_{cmux}^2 the variance added by an extended CMux using a GGSW coming from a circuit bootstrapping.

4 Addition and Subtraction of HFP

In this section, we detail the algorithms used to perform addition and subtraction operations with our floating-point representation. Initially, we describe the operations that manage the mantissa: the first aligns the two mantissas, and the second carries out the subtraction between them, followed by a realignment of the resulting value. Ultimately, the application of these algorithms enables us to efficiently implement homomorphic floating-point (HFP) addition and subtraction operations.

4.1 Managing Mantissas and Exponents

To add two floating-point numbers, we can not directly add their mantissas. First, we need their exponents to be equal and the mantissas to be aligned properly. In what follows, we describe the algorithms to homomorphically perform these operations.

4.1.1 Aligned Mantissa

Algorithm 7 takes as input ciphertexts encrypting two mantissas and their corresponding exponents, and returns the largest exponent \mathbf{e}_{\max} along with both aligned mantissas. The first step of this operation is to perform a subtraction between the two exponents to obtain $d = |\mathbf{e}_1 - \mathbf{e}_2|$ and the sign of this difference. The sign then allows us to select the larger exponent and the mantissa that needs to be aligned. Finally, a tree of CMux, using the bits of d , aligns the selected mantissa by removing the d least significant ciphertexts from the mantissa associated with \mathbf{e}_{\min} . All the steps of this operation are illustrated in Figure 2.

Lemma 7 (Aligned mantissa (Algorithm 7)) *Let $\mathbf{ct}_{\mathbf{m}_i}$ and $\mathbf{ct}_{\mathbf{e}_i}$ such that $\mathbf{ct}_{\mathbf{e}_i} = [\mathbf{ct}_{\mathbf{e}_i, \ell_{\mathbf{e}}-1}, \dots, \mathbf{ct}_{\mathbf{e}_i, 0}] \in [\text{LWE}_s(\mathbf{e}_{i, \ell_{\mathbf{e}}-1}), \dots, \text{LWE}_s(\mathbf{e}_{i, 0})]$ and $\mathbf{ct}_{\mathbf{m}_i} = [\mathbf{ct}_{\mathbf{m}_i, \ell_{\mathbf{m}}-1}, \dots, \mathbf{ct}_{\mathbf{m}_i, 0}] \in [\text{LWE}_s(\mathbf{m}_{i, \ell_{\mathbf{m}}-1}), \dots, \text{LWE}_s(\mathbf{m}_{i, 0})]$ with $i \in \{1, 2\}$ be two ciphertexts encrypting $\mathbf{m}_i \cdot (2^{\rho_{\mathbf{m}}})^{\mathbf{e}_i - \text{bias}}$. Let $(\mathbf{ct}_{\mathbf{m}'_{1\text{res}}}, \mathbf{ct}_{\mathbf{m}'_{2\text{res}}}, \mathbf{ct}_{\mathbf{e}_{\max}}) \leftarrow \text{AlignMantissa}(\mathbf{ct}_{\mathbf{e}_1}, \mathbf{ct}_{\mathbf{m}_1}, \mathbf{ct}_{\mathbf{e}_2}, \mathbf{ct}_{\mathbf{m}_2}, \text{PUB})$. Then, $\mathbf{e}_{\max} = \max(\mathbf{e}_1, \mathbf{e}_2)$. If $\mathbf{e}_1 > \mathbf{e}_2$, $\mathbf{m}'_{1\text{res}} = \mathbf{m}_1$, then $\mathbf{m}'_{2\text{res}} = \lfloor \mathbf{m}_2 / 2^{\rho_{\mathbf{m}} \cdot d} \rfloor$ with $d = \mathbf{e}_1 - \mathbf{e}_2$. Else if $\mathbf{e}_1 < \mathbf{e}_2$, $\mathbf{m}'_{2\text{res}} = \mathbf{m}_2$, then $\mathbf{m}'_{1\text{res}} = \lfloor \mathbf{m}_1 / 2^{\rho_{\mathbf{m}} \cdot d} \rfloor$ with $d = \mathbf{e}_2 - \mathbf{e}_1$. Else if $\mathbf{e}_1 = \mathbf{e}_2$, then $\mathbf{m}'_{1\text{res}} = \mathbf{m}_1$ and $\mathbf{m}'_{2\text{res}} = \mathbf{m}_2$. The complexity of the algorithm is: $\mathbb{C}_{\text{AlignMantissa}} = (\ell_{\mathbf{e}} \cdot \rho_{\mathbf{e}} + 1) \cdot \mathbb{C}_{\text{CBS}} + \mathbb{C}_{\text{IntSub}}^{\ell_{\mathbf{e}}}$*

Algorithm 7: $(\mathbf{ct}_{m_1'}, \mathbf{ct}_{m_2'}, \mathbf{ct}_e) \leftarrow \text{AlignMantissa}(\mathbf{ct}_{e_1}, \mathbf{ct}_{m_1}, \mathbf{ct}_{e_2}, \mathbf{ct}_{m_2}, \text{PUB})$

Input: $\begin{cases} \mathbf{ct}_{e_1} = [\mathbf{ct}_{e_1, \ell_e - 1}, \dots, \mathbf{ct}_{e_1, 0}] \in [\text{LWE}_s(\mathbf{e}_{1, \ell_e - 1}), \dots, \text{LWE}_s(\mathbf{e}_{1, 0})] \\ \mathbf{ct}_{m_1} = [\mathbf{ct}_{m_1, \ell_m - 1}, \dots, \mathbf{ct}_{m_1, 0}] \in [\text{LWE}_s(\mathbf{m}_{1, \ell_m - 1}), \dots, \text{LWE}_s(\mathbf{m}_{1, 0})] \\ \mathbf{ct}_{e_2} = [\mathbf{ct}_{e_2, \ell_e - 1}, \dots, \mathbf{ct}_{e_2, 0}] \in [\text{LWE}_s(\mathbf{e}_{2, \ell_e - 1}), \dots, \text{LWE}_s(\mathbf{e}_{2, 0})] \\ \mathbf{ct}_{m_2} = [\mathbf{ct}_{m_2, \ell_m - 1}, \dots, \mathbf{ct}_{m_2, 0}] \in [\text{LWE}_s(\mathbf{m}_{2, \ell_m - 1}), \dots, \text{LWE}_s(\mathbf{m}_{2, 0})] \end{cases}$
 PUB : Public materials for **PBS-KS** and for **CBS**

Output: $\begin{cases} \mathbf{ct}_{m_1'} = [\mathbf{ct}_{m_1', \ell_m - 1}, \dots, \mathbf{ct}_{m_1', 0}] \in [\text{LWE}_s(\mathbf{m}'_{1, \ell_m - 1}), \dots, \text{LWE}_s(\mathbf{m}'_{1, 0})] \\ \mathbf{ct}_{m_2'} = [\mathbf{ct}_{m_2', \ell_m - 1}, \dots, \mathbf{ct}_{m_2', 0}] \in [\text{LWE}_s(\mathbf{m}'_{2, \ell_m - 1}), \dots, \text{LWE}_s(\mathbf{m}'_{2, 0})] \\ \mathbf{ct}_e = [\mathbf{ct}_{e, \ell_e - 1}, \dots, \mathbf{ct}_{e, 0}] \in [\text{LWE}_s(\mathbf{e}_{\ell_e - 1}), \dots, \text{LWE}_s(\mathbf{e}_0)] \end{cases}$

/* Subtraction between the two exponents follows by the bit extraction */

```

1   $(\mathbf{ct}_s, \mathbf{ct}_d = [\mathbf{ct}_{d, \ell_e - 1}, \dots, \mathbf{ct}_{d, 0}]) \leftarrow \text{IntSub}^*(\mathbf{ct}_{e_1}, \mathbf{ct}_{e_2}, \text{PUB})$  ; /* Algo 2 */
2  for  $i$  in  $[0.. \ell_e - 1]$  do
3      for  $j$  in  $[0.. \rho_e - 1]$  do
4          /* Extract each bit of  $\mathbf{ct}_d$  */
5           $\overline{\overline{\text{CT}}}_{d(i \cdot \rho_e + j)} \leftarrow \text{CBS}(\mathbf{ct}_{d, i}, \text{PUB})$ 
6   $\overline{\overline{\text{CT}}}_{d_s} \leftarrow \text{CBS}(\mathbf{ct}_s, \text{PUB})$ 
7  /* selects the CT we need to align */
8   $\mathbf{ct}_{m_{\text{out}}}^{(0)} = [\mathbf{ct}_{m_{\text{out}}, \ell_m - 1}, \dots, \mathbf{ct}_{m_{\text{out}}, 0}] \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{m_2}, \mathbf{ct}_{m_1}, \overline{\overline{\text{CT}}}_{d_s})$  ; /* Algo 5 */
9  for  $i$  in  $[1.. \ell_m]$  do
10     /* Remove the  $i^{\text{th}}$  less significant blocks and add  $i$  trivial Zero LWEs on the most significant blocks */
11      $\mathbf{ct}_{m_{\text{out}}}^{(i)} \leftarrow [\text{TrivialEncrypt}(0, i), \mathbf{ct}_{m_{\text{out}}, \ell_m - 1}, \dots, \mathbf{ct}_{m_{\text{out}}, i}]$ 
12     for  $i$  in  $[0.. \ell_e \cdot \rho_e]$  do
13         if  $\lfloor \ell_m / 2^{i+1} \rfloor = 0$  then
14              $\mathbf{ct}_{m_{\text{out}}}^{(0)} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{m_{\text{out}}}^{(0)}, \text{TrivialEncrypt}(0, \ell_m), \overline{\overline{\text{CT}}}_{d_i})$  ; /* Algo 5 */
15         else
16             for  $j$  in  $[0.. \lfloor \ell_m / 2^{i+1} \rfloor]$  do
17                 /* If  $\mathbf{ct}_{m_{\text{out}}}^{(2 \cdot j + 1)}$  is not defined, then it is equal to  $\text{TrivialEncrypt}(0, \ell_m)$  */
18                  $\mathbf{ct}_{m_{\text{out}}}^{(i)} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{m_{\text{out}}}^{(2 \cdot j)}, \mathbf{ct}_{m_{\text{out}}}^{(2 \cdot j + 1)}, \overline{\overline{\text{CT}}}_{d_i})$  ; /* Algo 5 */
19   $\mathbf{ct}_{m_1'} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{m_1}, \mathbf{ct}_{m_{\text{out}}}^{(0)}, \overline{\overline{\text{CT}}}_{d_s})$ ,
20   $\mathbf{ct}_{m_2'} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{m_{\text{out}}}^{(0)}, \mathbf{ct}_{m_2}, \overline{\overline{\text{CT}}}_{d_s})$  ; /* Algo 5 */
21   $\mathbf{ct}_e \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{e_1}, \mathbf{ct}_{e_2}, \overline{\overline{\text{CT}}}_{d_s})$  ; /* Algo 5 */
22  return  $(\mathbf{ct}_{m_1'}, \mathbf{ct}_{m_2'}, \mathbf{ct}_e)$ 
    
```

Proof 7 (Correctness of AlignMantissa (Algorithm 7)) *In what follows, we use the index of a ciphertext to refer to the plaintext value it encrypts, i.e., $\mathfrak{z} = \text{Decrypt}_{\mathfrak{s}}(\text{ct}_{\mathfrak{z}})$. From Algorithm 2, we have that $d = |\mathfrak{e}_1 - \mathfrak{e}_2|$ and $\mathfrak{s} = 0$ if $\mathfrak{e}_1 \geq \mathfrak{e}_2$, 1 otherwise. Each bits of $\text{ct}_{\mathfrak{s}}$ and $\text{ct}_{\mathfrak{d}}$ are converted to GGSW via a CBS, so that after Line 5, we have: $\{\overline{\text{CT}}_{d_i} \in \text{GGSW}_{\mathfrak{s}}\}_{(d_i)_{i \in [0, \ell_{\mathfrak{e}} \cdot \rho_{\mathfrak{e}} - 1]}}$ (such that $d = \sum_{i=0}^{\ell_{\mathfrak{e}} \cdot \rho_{\mathfrak{e}}} d_i 2^i$) and $\overline{\text{CT}}_{\mathfrak{s}} \in \text{GGSW}_{\mathfrak{s}}(d_{\mathfrak{s}})$. From Algorithm 5, after Line 6, we get $\mathbf{m}_{\text{out}}^{(0)} = \mathbf{m}_2$ if $d_{\mathfrak{s}} = 0$, \mathbf{m}_1 otherwise. This gives which of the mantissa needs to be aligned, i.e., if the ciphertext $\overline{\text{CT}}_{d_{\mathfrak{s}}}$ is in $\text{GGSW}_{\mathfrak{s}}^{\beta, \ell}(0)$, $\mathfrak{e}_1 \geq \mathfrak{e}_2$ so we need to align \mathbf{m}_2 , otherwise $\mathfrak{e}_1 < \mathfrak{e}_2$, and thus \mathbf{m}_1 needs to be aligned). In next loop, encryption of all possible mantissa shifts are created, i.e., $M^{(0)} = \{\mathbf{ct}_{\text{mout}}^{(i)} = [\mathbf{0}, \dots, \mathbf{0}, \text{ct}_{\text{mout}, \ell_{\mathfrak{m}} - 1}, \dots, \text{ct}_{\text{mout}, i}]\}_{i \in [1, \ell_{\mathfrak{m}} - 1]}$ with $\mathbf{0} \in \text{LWE}_{\mathfrak{s}}(0)$, s.t. $\mathbf{m}_{\text{out}}^{(i)} = \lfloor \mathbf{m}_{\text{out}}^{(0)} / 2^{\rho_{\mathfrak{m}} \cdot i} \rfloor$. The next steps consists in choosing the right mantissa from this set, depending on the value of d . Informally, d is the number of blocks by which the mantissa should be shifted. At each step i of the loop, the set $M^{(i)}$ is updated with respect of the binary value of d , to contains the encryption of each value in $\{\mathbf{ct}_{\text{mout}}^{(\alpha)} = \lfloor \mathbf{m}_{\text{out}}^{(0)} / 2^{\rho_{\mathfrak{m}} \cdot \alpha} \rfloor \text{ s.t. } \alpha = \sum_{j=0}^{i-1} d_j 2^j \bmod 2^{i+1}\}$. Note that in the case where $\mathfrak{e}_1 = \mathfrak{e}_2$, the cmux tree will return the selected mantissa without any change. In the end, the set is reduced to a singleton containing $\lfloor \mathbf{m}_{\text{out}}^{(0)} / 2^{\rho_{\mathfrak{m}} \cdot d} \rfloor$. Finally, the last three CMux replace the mantissa value by the aligned one and select the bigger exponent.*

Lemma 8 (Noise Constraints of Algorithm 7) *The output noise variances of the ciphertexts of Algorithm 7, $\text{ct}_{\mathfrak{m}_i'}$ and $\text{ct}_{\mathfrak{e}}$, are respectively $\sigma_{\text{in}, \mathfrak{m}}^2 + \frac{(\rho_{\mathfrak{e}} \cdot \ell_{\mathfrak{e}} + 3)}{2} \cdot \sigma_{\text{cmux}}^2$ and $\sigma_{\text{in}, \mathfrak{e}}^2 + \sigma_{\text{cmux}}^2$.*

To guarantee correctness of Algorithm 7, we need to find parameters that verify the following inequalities:

$$4 \cdot \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$$

With $\sigma_{\text{in}, \mathfrak{e}}$ the noise variance of the input exponent ciphertexts, $\sigma_{\text{in}, \mathfrak{m}}$ the noise variance of the input mantissa ciphertexts and with σ_{BR} the noise added by the blind rotation, σ_{KS} the noise added by the keyswitch, σ_{CMux} the noise added by the CMux and finally σ_{MS} , the noise added by the modulus switch.

Proof 8 (Proof of Lemma 8) *Let us assume that the input ciphertexts $\text{ct}_{\mathfrak{e}_i}$ and $\text{ct}_{\mathfrak{m}_i}$ have respectively a noise variance $\sigma_{\text{in}, \mathfrak{e}}^2$ and $\sigma_{\text{in}, \mathfrak{m}}^2$. The first*

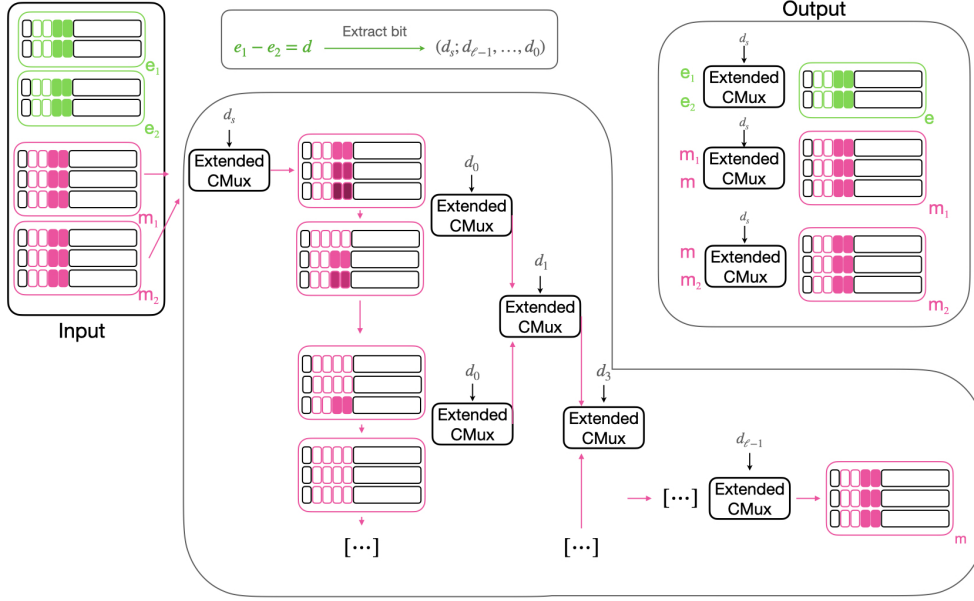


Figure 2: This figure illustrates the main steps of the Homomorphic Align-Mantissa operation (see Algorithm 7). The goal is to align the mantissas based on their exponents. At a high level, we first use Algorithm 2 on the two exponents (in green) to compute the difference between the two exponents, along with the sign. Using the sign, we can determine which mantissa (in pink) is smaller. Then, with the difference and a tree of Extended CMux operations, we can decide how many ciphertexts are needed to increase the smaller mantissa and align the exponents. The final step involves correctly ordering the two mantissas and selecting the larger exponent.

line calls Algorithm 2 (see Lemma 2 for more details). In particular, we compute the output noise of Algorithm 2. The noise variances of \mathbf{ct}_s and \mathbf{ct}_d are respectively $4 \cdot \sigma_{\text{BR}}^2$ and σ_{BR}^2 with σ_{BR}^2 , the noise variance of a freshly bootstrapped ciphertext. In the next lines, the algorithm heavily relies on circuit bootstrapping which gives us the following noise constraints : $\sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$ & $4 \cdot \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$, with σ_{KS}^2 and σ_{MS}^2 respectively the noise variance added by a keyswitch and by a modulus switch. t^2 represents the noise bound as previously defined in the proof of Lemma 6. As the first constraint is dominated by the second, we can remove it from

the set of constraints. Then, we apply an extended cmux and we create ℓ_m vectors of ciphertexts composed of outputs of the previous extended cmux and trivial ciphertexts. At this stage, we assume that every ciphertext in these vectors has the same noise variance $\sigma_{\text{in},m}^2 + \sigma_{\text{cmux}}^2$ with σ_{cmux}^2 the noise added by a cmux using a GGSW coming from a circuit bootstrap with a noise variance $\sigma_{\text{BR}}^2 + \sigma_{\text{PPKS}}^2$. Then, we apply an extended cmux tree of depth $\rho_e \cdot \ell_e$. Therefore, at the end of the tree, the ciphertexts have a noise variance $\sigma_{\text{in},m}^2 + (\rho_e \cdot \ell_e + 1) \cdot \sigma_{\text{cmux}}^2$. At the end of the algorithm, assuming that \mathbf{ct}_s encrypt 0 or 1 with probability $\frac{1}{2}$, the noise variance of $\mathbf{ct}_{m'_i}$ is $\sigma_{\text{in},m}^2 + \frac{(\rho_e \cdot \ell_e + 3)}{2} \cdot \sigma_{\text{cmux}}^2$. The noise variance of \mathbf{ct}_e is $\sigma_{\text{in},e}^2 + \sigma_{\text{cmux}}^2$.

4.1.2 SubMantissa

SubMantissa performs the subtraction of two mantissas and shifts the result such that the most significant block is not empty (unless the result is zero or too small to be represented). It changes the value of the exponent and the sign consequently. To perform this operation, the two mantissas must to be aligned. Algorithm 8 takes as input the encryption of two aligned mantissas, the exponent and the sign, and returns the encryption of the absolute value of the difference of the mantissas, the exponent and the sign of this subtraction.

Lemma 9 (SubMantissa (Algorithm 8)) *Let $\mathbf{ct}_{m_i} = [\mathbf{ct}_{m_i, \ell_m - 1}, \dots, \mathbf{ct}_{m_i, 0}] \in [\text{LWE}_s(\mathbf{m}_{i, \ell_m - 1}), \dots, \text{LWE}_s(\mathbf{m}_{i, 0})]$ with $i \in \{0, 1\}$ be ciphertexts encrypting $\mathbf{m}_i < 2^{\rho_m}$. Let $\mathbf{ct}_e = [\mathbf{ct}_{e, \ell_e - 1}, \dots, \mathbf{ct}_{e, 0}] \in [\text{LWE}_s(\mathbf{e}_{\ell_e - 1}), \dots, \text{LWE}_s(\mathbf{e}_0)]$ a ciphertexts encrypting $\mathbf{e} < 2^{\rho_e}$. Let $\mathbf{ct}_{s_1} \in \text{LWE}_s(\mathbf{s}_1)$ and $\mathbf{ct}_{s_2} \in \text{LWE}_s(\mathbf{s}_2)$ with $\mathbf{s}_1 = 1 - \mathbf{s}_2$ such that $\mathbf{f}_1 = (-1)^{\mathbf{s}_1} \cdot \mathbf{m}_1 \cdot (2^{\rho_m})^{\mathbf{e} - \text{bias}}$ and $\mathbf{f}_2 = (-1)^{\mathbf{s}_2} \cdot \mathbf{m}_2 \cdot (2^{\rho_m})^{\mathbf{e} - \text{bias}}$. Let $\mathbf{ct}_{\mathbf{f}_{\text{sub}}} = (\mathbf{ct}_{\mathbf{m}_{\text{sub}}}, \mathbf{ct}_{\mathbf{e}_{\text{sub}}}, \mathbf{ct}_{\mathbf{s}_{\text{sub}}}) \leftarrow \text{SubMantissa}(\mathbf{ct}_{m_{\text{in},1}}, \mathbf{ct}_{m_{\text{in},2}}, \mathbf{ct}_e, \mathbf{ct}_{s_1}, \text{PUB})$. Then $\text{DecryptFloat}_s(\mathbf{ct}_{\mathbf{f}_{\text{sub}}}) = \mathbf{f}_1 - \mathbf{f}_2 = (-1)^{\mathbf{s}_{\text{sub}}} \cdot \mathbf{m}_{\text{sub}} \cdot (2^{\rho_m})^{\mathbf{e}_{\text{sub}} - \text{bias}}$ such that $\mathbf{s}_{\text{sub}} = \mathbf{s}_1$ if $\mathbf{m}_1 \geq \mathbf{m}_2$, or \mathbf{s}_2 if $\mathbf{m}_1 < \mathbf{m}_2$. Assuming $\mathbf{m}_1 \neq \mathbf{m}_2$, let α be the index of the first non zero block of $\mathbf{m} = |\mathbf{m}_1 - \mathbf{m}_2|$ i.e., $\alpha = \min_{i \in [0, \ell_m - 1]} \{\ell_m - 1 - i \text{ s.t. } \mathbf{m}_i \neq 0\}$ then, if $\mathbf{e} \geq \alpha$ $\mathbf{e}_{\text{sub}} = \mathbf{e} - \alpha$ and $\mathbf{m}_{\text{sub}} = |\mathbf{m}_1 - \mathbf{m}_2| \cdot 2^{\rho_m \cdot \alpha}$. Else if $\mathbf{m}_1 = \mathbf{m}_2$ or if $\mathbf{e} - \alpha < 0$, then, $\mathbf{m}_{\text{sub}} = 0$, $\mathbf{e}_{\text{sub}} = 0$.*

The complexity of the algorithm is: $\mathbb{C}_{\text{SubMantissa}} = \mathbb{C}_{\text{IntSub}^}^{\ell_m} + \mathbb{C}_{\text{IntSub}^*}^{\ell_e} + (\ell_m + 1) \cdot \mathbb{C}_{\text{CBS}}$*

Algorithm 8: $(\mathbf{ct}_{\mathbf{m}_{\text{sub}}}, \mathbf{ct}_{\mathbf{e}_{\text{sub}}}, \mathbf{ct}_{\mathbf{s}_{\text{sub}}}) \leftarrow \text{SubMantissa}(\mathbf{ct}_{\mathbf{m}_{\text{in},1}}, \mathbf{ct}_{\mathbf{m}_{\text{in},2}}, \mathbf{ct}_{\mathbf{e}}, \mathbf{ct}_{\mathbf{s}_1}, \text{PUB})$

Input: $\begin{cases} \mathbf{ct}_{\mathbf{m}_{\text{in},1}} = [\mathbf{ct}_{\mathbf{m}_{\text{in},1}, \ell_{\mathbf{m}}-1}, \dots, \mathbf{ct}_{\mathbf{m}_{\text{in},1}, 0}] \in [\text{LWE}_{\mathbf{s}}(\mathbf{m}_{\text{in},1, \ell_{\mathbf{m}}-1}), \dots, \text{LWE}_{\mathbf{s}}(\mathbf{m}_{\text{in},1, 0})] \\ \mathbf{ct}_{\mathbf{m}_{\text{in},2}} = [\mathbf{ct}_{\mathbf{m}_{\text{in},2}, \ell_{\mathbf{m}}-1}, \dots, \mathbf{ct}_{\mathbf{m}_{\text{in},2}, 0}] \in [\text{LWE}_{\mathbf{s}}(\mathbf{m}_{\text{in},2, \ell_{\mathbf{m}}-1}), \dots, \text{LWE}_{\mathbf{s}}(\mathbf{m}_{\text{in},2, 0})] \\ \mathbf{ct}_{\mathbf{e}} = [\mathbf{ct}_{\mathbf{e}, \ell_{\mathbf{e}}-1}, \dots, \mathbf{ct}_{\mathbf{e}, 0}] \in [\text{LWE}_{\mathbf{s}}(\mathbf{e}_{\ell_{\mathbf{e}}-1}), \dots, \text{LWE}_{\mathbf{s}}(\mathbf{e}_0)] \\ \mathbf{ct}_{\mathbf{s}_1} \in \text{LWE}_{\mathbf{s}}(\text{sign}_1) \\ \text{PUB} : \text{Public keys for KS-PBS and for CBS} \end{cases}$

Output: $\begin{cases} \mathbf{ct}_{\mathbf{m}_{\text{sub}}} = [\mathbf{ct}_{\mathbf{m}_{\text{sub}}, \ell_{\mathbf{m}}-1}, \dots, \mathbf{ct}_{\mathbf{m}_{\text{sub}}, 0}] \in [\text{LWE}_{\mathbf{s}}(\mathbf{m}_{\text{sub}, \ell_{\mathbf{m}}-1}), \dots, \text{LWE}_{\mathbf{s}}(\mathbf{m}_{\text{sub}, 0})] \\ \mathbf{ct}_{\mathbf{e}_{\text{sub}}} = [\mathbf{ct}_{\mathbf{e}_{\text{sub}}, \ell_{\mathbf{e}}-1}, \dots, \mathbf{ct}_{\mathbf{e}_{\text{sub}}, 0}] \in [\text{LWE}_{\mathbf{s}}(\mathbf{e}_{\text{sub}, \ell_{\mathbf{e}}-1}), \dots, \text{LWE}_{\mathbf{s}}(\mathbf{e}_{\text{sub}, 0})] \\ \mathbf{ct}_{\mathbf{s}_{\text{sub}}} \in \text{LWE}_{\mathbf{s}}(\mathbf{s}_{\text{sub}}) \end{cases}$

- 1 $(\mathbf{ct}_{\mathbf{s}}, \mathbf{ct}_{\mathbf{m}_0}) = [\mathbf{ct}_{\mathbf{m}_0, \ell_{\mathbf{m}}-1}, \dots, \mathbf{ct}_{\mathbf{m}_0, 0}] \leftarrow \text{IntSub}^*(\mathbf{ct}_{\mathbf{m}_{\text{in},1}}, \mathbf{ct}_{\mathbf{m}_{\text{in},2}}, \text{PUB})$
- 2 $\mathbf{ct}_{\mathbf{e}_0} \leftarrow \text{TrivialEncrypt}(0, \ell_{\mathbf{e}})$
- 3 **for** i **in** $[1.. \ell_{\mathbf{m}}]$ **do**
- 4 $\mathbf{ct}'_{\mathbf{e}_i} \leftarrow \text{TrivialEncrypt}(i, \ell_{\mathbf{e}})$
- 5 $\mathbf{ct}_0 \leftarrow \text{TrivialEncrypt}(0, 1)$
- 6 $\mathbf{ct}_{\mathbf{m}_i} \leftarrow [\mathbf{ct}_{\mathbf{m}_{i-1}, \ell_{\mathbf{m}}-2}, \dots, \mathbf{ct}_{\mathbf{m}_{i-1}, 0}, \mathbf{ct}_0]$
- 7 $\overline{\overline{\mathbf{CT}}} \in \text{GGSW}_{\mathbf{S}}^{\beta, \ell}(0)$ **if** $\mathbf{ct}_{\mathbf{m}_{i-1}, \ell_{\mathbf{m}}-1} \in \text{LWE}_{\mathbf{s}}(0)$, $\overline{\overline{\mathbf{CT}}} \in \text{GGSW}_{\mathbf{S}}^{\beta, \ell}(1)$ **otherwise** */
- 8 $\overline{\overline{\mathbf{CT}}} \leftarrow \text{CBS}(\mathbf{ct}_{\mathbf{m}_{i-1}, \ell_{\mathbf{m}}-1}, \text{PUB})$
- 9 $\mathbf{ct}_{\mathbf{m}_i} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{\mathbf{m}_i}, \mathbf{ct}_{\mathbf{m}_{i-1}}, \overline{\overline{\mathbf{CT}}})$
- 10 **if** $i \neq \ell_{\mathbf{m}}$ **then** $\mathbf{ct}'_{\mathbf{e}_i} \leftarrow \text{ExtendedCMux}(\mathbf{ct}'_{\mathbf{e}_i}, \mathbf{ct}'_{\mathbf{e}_{i-1}}, \overline{\overline{\mathbf{CT}}})$ **else**
- 11 $\mathbf{ct}'_{\mathbf{e}_i} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{\mathbf{e}}, \mathbf{ct}'_{\mathbf{e}_{i-1}}, \overline{\overline{\mathbf{CT}}})$;
- 12 $(\mathbf{ct}_{\mathbf{s}_{\mathbf{e}}}, \mathbf{ct}_{\mathbf{e}_{\text{res}}}) \leftarrow \text{IntSub}^*(\mathbf{ct}_{\mathbf{e}}, \mathbf{ct}'_{\mathbf{e}_{\ell_{\mathbf{m}}}}, \text{PUB})$; /* Algo 2 */
- 13 $\overline{\overline{\mathbf{CT}}}_{\mathbf{s}_{\mathbf{e}}} \in \text{GGSW}_{\mathbf{S}}^{\beta, \ell}(0)$ **if** $\mathbf{ct}_{\mathbf{s}_{\mathbf{e}}} \in \text{LWE}_{\mathbf{s}}(0)$, $\overline{\overline{\mathbf{CT}}}_{\mathbf{s}_{\mathbf{e}}} \in \text{GGSW}_{\mathbf{S}}^{\beta, \ell}(1)$ **otherwise** */
- 14 $\overline{\overline{\mathbf{CT}}}_{\mathbf{s}_{\mathbf{e}}} \leftarrow \text{CBS}(\mathbf{ct}_{\mathbf{s}_{\mathbf{e}}}, \text{PUB})$
- 15 $\mathbf{ct}_{\mathbf{m}_{\text{sub}}} \leftarrow \text{ExtendedCMux}(\text{TrivialEncrypt}(0, \ell_{\mathbf{m}}), \mathbf{ct}_{\mathbf{m}_{\ell_{\mathbf{m}}}}, \overline{\overline{\mathbf{CT}}}_{\mathbf{s}_{\mathbf{e}}})$,
- 16 $\mathbf{ct}_{\mathbf{e}_{\text{sub}}} \leftarrow \text{ExtendedCMux}(\text{TrivialEncrypt}(0, \ell_{\mathbf{e}}), \mathbf{ct}_{\mathbf{e}_{\text{res}}}, \overline{\overline{\mathbf{CT}}}_{\mathbf{s}_{\mathbf{e}}})$
- 17 $\mathbf{ct}_{\mathbf{s}_{\text{sub}}} \leftarrow \mathbf{ct}_{\mathbf{s}_1} + \mathbf{ct}_{\mathbf{s}}$
- 18 **return** $(\mathbf{ct}_{\mathbf{e}_{\text{sub}}}; \mathbf{ct}_{\mathbf{m}_{\text{sub}}}; \mathbf{ct}_{\mathbf{s}_{\text{sub}}})$

Proof 9 (Correctness of SubMantissa (Algorithm 8)) *The first step of the algorithm is to subtract the two mantissas. We obtain $\mathbf{ct}_{\mathbf{m}_0}$ which is equals to $|\mathbf{ct}_{\mathbf{m}_{\text{in},1}} - \mathbf{ct}_{\mathbf{m}_{\text{in},2}}|$ and $\mathbf{ct}_{\mathbf{s}}$ the sign of this subtraction. As the two mantissas are aligned, we have \mathbf{m}_0 in $[0, 2^{\ell_{\mathbf{m}} \cdot \rho_{\mathbf{m}}})$.*

At each step i of the loop, we take the previously computed ciphertext $\text{ct}_{m_{i-1}}$ encrypting a message m_{i-1} and build another ciphertext ct_{m_i} encrypting the message $m_i = m_{i-1} \cdot 2^{\rho_m}$. On line 7, we create a GGSW ciphertext $\overline{\overline{\text{CT}}}$ encrypting 0 if the most significant block of the mantissa m_{i-1, ℓ_m-1} is equal to 0 and 1 if it contains some non-zero integer. Remember, our goal is to realign the mantissa to stay in the classical representation (i.e., $m_{\text{sub}} \in [2^{(\ell_m-1) \cdot \rho_m}, 2^{\ell_m \cdot \rho_m})$ or $m_{\text{sub}} = 0$ if $m_{\text{in},1} = m_{\text{in},2}$). Therefore, we use a cmux to select $\text{ct}_{m_{i-1}}$ if $m_{i-1, \ell_m-1} \neq 0$ and ct_{m_i} if $m_{i-1, \ell_m-1} = 0$. In the same way, we use the cmux to update the value of the exponent. To do so, we take the previously computed $\text{ct}_{e_{i-1}}$ and create a new ciphertext ct_{e_i} , a trivial encryption of i . As we want to select $\text{ct}_{e_{i-1}}$ (respectively ct_{e_i}) if we selected $\text{ct}_{m_{i-1}}$ (resp. ct_{m_i}) in the previous step, we can perform a cmux with the same GGSW ciphertext $\overline{\overline{\text{CT}}}$. Assuming that we have $m_0 \neq 0$, at the end of the for-loop, ct_{e_i} is encrypting a value α such that $m_0 \cdot 2^{\rho_m \cdot \alpha} \in [2^{(\ell_m-1) \cdot \rho_m}, 2^{\ell_m \cdot \rho_m})$ and ct_{m_i} is encrypting $m_0 \cdot 2^{\rho_m \cdot \alpha}$. If we have $m_0 = 0$, ct_{m_i} will still be equal to 0.

The next step is to update the exponent. In line 13, we subtract the value α encrypted in $\text{ct}_{e_{\ell_m}}$ to ct_e . The sign of this subtraction is in $\text{LWE}_s(0)$ if we can do the subtraction ($e > \alpha$) otherwise, the result is in $\text{LWE}_s(1)$, the value of the subtraction is too small to be represented with our encoding. With this sign we create a new GGSW ciphertext $\overline{\overline{\text{CT}}}$. Finally, the last CMux returns the ciphertexts encrypting mantissa $m_{\text{sub}} = m_0 \cdot 2^{\rho_m \cdot \alpha}$ and the ciphertexts encrypting the exponent $e_{\text{sub}} = e - \alpha$ if the subtraction can be done, otherwise it returns the ciphertexts encrypting 0.

Lemma 10 (Noise Constraints of Algorithm 8) *The output noise variances of ciphertexts of Algorithm 8, $\text{ct}_{m_{\text{sub}}}$, $\text{ct}_{e_{\text{sub}}}$ and $\text{ct}_{s_{\text{sub}}}$, are respectively $\sigma_{\text{BR}}^2 + (\ell_m + 1) \cdot \sigma_{\text{cmux}}^2$, $\sigma_{\text{BR}}^2 + \sigma_{\text{cmux}}^2$ and $\sigma_{\text{in},s}^2 + 4\sigma_{\text{BR}}^2$.*

To guarantee correctness of Algorithm 8, we need to find parameters that verify the following inequalities:

$$\sigma_{\text{BR}}^2 + (i-1) \cdot \sigma_{\text{cmux}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2 \quad \& \quad 4\sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$$

With $\sigma_{\text{in},s}$ the noise variance of the input sign ciphertext and with σ_{BR} the noise added by the blind rotation, σ_{KS} the noise added by the keyswitch, σ_{CMux} the noise added by the CMux and finally σ_{MS} , the noise added by the modulus switch.

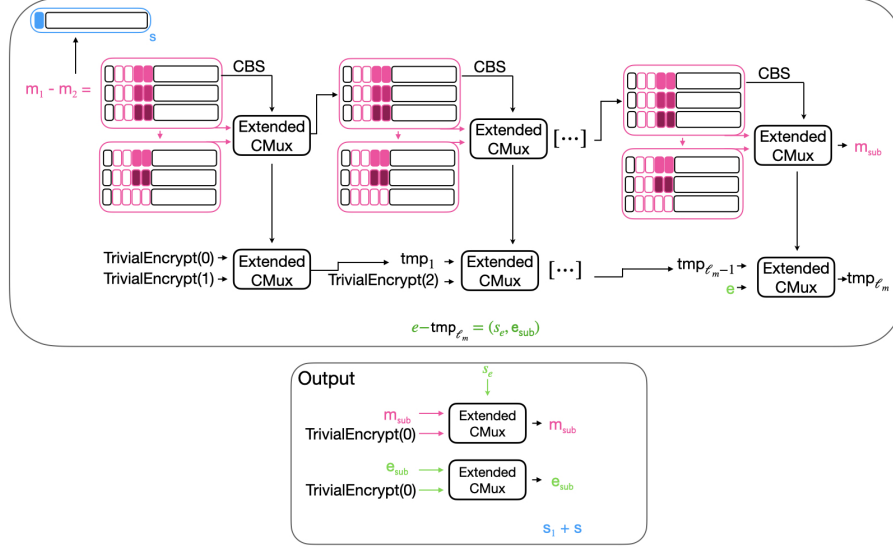


Figure 3: This figure represents the main steps of Algorithm 8. At a high level, we first subtract the two mantissas using Algorithm 2, obtaining the absolute value of the subtraction along with the sign. Next, we perform a loop where, at each step, the first ciphertext of the new mantissa is transformed into a GGSW ciphertext using a CBS. Using this GGSW ciphertext, we remove the most significant ciphertext if it is empty otherwise, we keep the mantissa. At the same time, at each step, we count the number of ciphertexts removed. When the loop finishes, we subtract the number of removed ciphertexts from the exponent. Finally, we use an extended Cmux to return 0 if the exponent is negative or if the mantissa is empty.

Proof 10 (Proof of Lemma 10) *Let us assume that the input ciphertexts \mathbf{ct}_e , \mathbf{ct}_{m_i} and \mathbf{ct}_{s_1} have respectively a noise variance $\sigma_{\text{in},e}^2$, $\sigma_{\text{in},m}^2$ and $\sigma_{\text{in},s}^2$.*

Using the noise analysis of Algorithm 2 presented in Lemma 2, we know that the noise variance of \mathbf{ct}_{m_0} and \mathbf{ct}_s are respectively σ_{BR}^2 and $4 \cdot \sigma_{\text{BR}}^2$. In the for-loop, for each index $1 \leq i \leq \ell_m$, we can compute the noise constraint $\sigma_{\text{BR}}^2 + (i-1) \cdot \sigma_{\text{cmux}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$ and the noise variance of \mathbf{ct}_{m_i} is $\sigma_{\text{BR}}^2 + i \cdot \sigma_{\text{cmux}}^2$, the noise variance of \mathbf{ct}_{e_i} is $i \cdot \sigma_{\text{cmux}}^2$ for $i \neq \ell_m$ and $\max((\ell_m - 1) \cdot \sigma_{\text{cmux}}^2, \sigma_{\text{in},e}^2) + \sigma_{\text{cmux}}^2$ for $i = \ell_m$. As in the previous proofs, we only need to retain the noise constraint for $i = \ell_m$, as it dominates the other constraints. Then, we

have a call to Algorithm 2 which gives us ciphertexts of variances respectively $4\sigma_{\text{BR}}^2$ and σ_{BR}^2 . Next, we perform a circuit bootstrapping which gives us the following constraint $4\sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$. Finally, the algorithm outputs $\text{ct}_{\text{m}_{\text{sub}}}, \text{ct}_{\text{e}_{\text{sub}}}$ and $\text{ct}_{\text{s}_{\text{sub}}}$ of respective variances $\sigma_{\text{BR}}^2 + (\ell_{\text{m}} + 1) \cdot \sigma_{\text{cmux}}^2$, $\sigma_{\text{BR}}^2 + \sigma_{\text{cmux}}^2$ and $\sigma_{\text{in},s}^2 + 4\sigma_{\text{BR}}^2$.

4.2 Addition and Subtraction

This operation performs the addition of two homomorphic floating-point numbers. To perform a subtraction, we only need to change the input sign of the second ciphertext. This operation is straightforward, as the sign is on a padding bit, adding the clear integer $q/2$ to the sign ciphertext change the sign of the floating-point.

This operation is based on the previous algorithm. We first need to align the mantissas (Algorithm 7). Next, we perform both the addition and the subtraction (Algorithm 8) of the mantissas and then we choose which of the two results to output based on the signs. All the steps of this operation are illustrated in Figure 4.

After this operation, we need to perform the operation CarryPropagateFloat (Algorithm 6) to retrieve the homomorphic floating-point representation.

Lemma 11 (Addition (Algorithm 9)) *Let ct_{f_i} such that $\text{ct}_{\text{s}_i} \in \text{LWE}_{\text{s}}(\text{s}_i), \text{ct}_{\text{e}_i} = [\text{ct}_{\text{e}_i, \ell_{\text{e}}-1}, \dots, \text{ct}_{\text{e}_i, 0}] \in [\text{LWE}_{\text{s}}(\text{e}_{i, \ell_{\text{e}}-1}), \dots, \text{LWE}_{\text{s}}(\text{e}_{i, 0})]$ encrypting $\text{e}_i < 2^{\rho_{\text{e}}}$ and $\text{ct}_{\text{m}_i} = [\text{ct}_{\text{m}_i, \ell_{\text{m}}-1}, \dots, \text{ct}_{\text{m}_i, 0}] \in [\text{LWE}_{\text{s}}(\text{m}_{i, \ell_{\text{m}}-1}), \dots, \text{LWE}_{\text{s}}(\text{m}_{i, 0})]$ encrypting $\text{m}_i < 2^{\rho_{\text{m}}}$ with $i \in \{1, 2\}$ be two ciphertexts encrypting $\text{f}_i = (-1)^{\text{s}_i} \cdot \text{m}_i \cdot (2^{\rho_{\text{m}}})^{\text{e}_i - \text{bias}}$. Let $(\text{ct}_{\text{m}_{\text{res}}}, \text{ct}_{\text{e}_{\text{res}}}, \text{ct}_{\text{s}_{\text{res}}}) = \text{ct}_{\text{f}_{\text{res}}} \leftarrow \text{Addition}(\text{ct}_{\text{f}_1}, \text{ct}_{\text{f}_2}, \text{PUB})$. Then $\text{DecryptFloat}(\text{ct}_{\text{f}_{\text{res}}}) = (-1)^{\text{s}_{\text{res}}} \cdot \text{m}_{\text{res}} \cdot (2^{\rho_{\text{m}}})^{\text{e}_{\text{res}} - \text{bias}} = \text{f}_{\text{res}} = \text{f}_1 + \text{f}_2 + \epsilon$ such that if $\text{e}_1 \geq \text{e}_2$, then $\text{m}'_1 = \text{m}_1$ and $\text{m}'_2 = \lfloor \text{m}_2 / 2^{\rho_{\text{m}} \cdot \gamma} \rfloor$ with $\gamma = \text{e}_1 - \text{e}_2$. Else if $\text{e}_1 \leq \text{e}_2$, then $\text{m}'_2 = \text{m}_2$ and $\text{m}'_1 = \lfloor \text{m}_1 / 2^{\rho_{\text{m}} \cdot \gamma} \rfloor$ with $\gamma = \text{e}_2 - \text{e}_1$. For ϵ , we refer to Definition 5. If $\text{s}_1 = \text{s}_2$, then $\text{s}_{\text{res}} = \text{s}_1$, $\text{e}_{\text{res}} = \max(\text{e}_1, \text{e}_2)$ and $\text{m}_{\text{res}} = \text{m}'_1 + \text{m}'_2$. Else if $\text{s}_1 \neq \text{s}_2$ and if $\text{m}'_1 \geq \text{m}'_2$, then $\text{s}_{\text{res}} = \text{s}_1$; or if $\text{m}'_1 < \text{m}'_2$, then $\text{s}_{\text{res}} = \text{s}_2$. Assuming $\text{m}'_1 \neq \text{m}'_2$ and $\text{s}_1 \neq \text{s}_2$, let α be the index of the first non zero block of $\text{m} = |\text{m}'_1 - \text{m}'_2|$, if $\max(\text{e}_1, \text{e}_2) \geq \alpha$, then $\text{e}_{\text{res}} = \max(\text{e}_1, \text{e}_2) - \alpha$ and $\text{m}_{\text{res}} = |\text{m}'_1 - \text{m}'_2| \cdot 2^{\ell_{\text{m}} \cdot \alpha}$. Else if $\max(\text{e}_1, \text{e}_2) < \alpha$, then $\text{e}_{\text{res}} = 0$, $\text{m}_{\text{res}} = 0$ and $\text{e}_{\text{res}} = 0$. The complexity of the algorithm is: $\mathbb{C}_{\text{Addition}} = \mathbb{C}_{\text{SubMantissa}} + \mathbb{C}_{\text{AlignMantissa}} + \mathbb{C}_{\text{CBS}}$.*

Algorithm 9: $\mathbf{ct}_{f_{res}} \leftarrow \text{Addition}(\mathbf{ct}_{f_1}, \mathbf{ct}_{f_2}, \text{PUB})$

Input: $\begin{cases} \mathbf{ct}_{s_1} \in \text{LWE}_s(s_1) \\ \mathbf{ct}_{e_1} = [\mathbf{ct}_{e_{1,\ell_e-1}}, \dots, \mathbf{ct}_{e_{1,0}}] \in [\text{LWE}_s(e_{1,\ell_e-1}), \dots, \text{LWE}_s(e_{1,0})] \\ \mathbf{ct}_{m_1} = [\mathbf{ct}_{m_{1,\ell_m-1}}, \dots, \mathbf{ct}_{m_{1,0}}] \in [\text{LWE}_s(m_{1,\ell_m-1}), \dots, \text{LWE}_s(m_{1,0})] \end{cases}$
 $\begin{cases} \mathbf{ct}_{s_2} \in \text{LWE}_s(s_2) \\ \mathbf{ct}_{e_2} = [\mathbf{ct}_{e_{2,\ell_e-1}}, \dots, \mathbf{ct}_{e_{2,0}}] \in [\text{LWE}_s(e_{2,\ell_e-1}), \dots, \text{LWE}_s(e_{2,0})] \\ \mathbf{ct}_{m_2} = [\mathbf{ct}_{m_{2,\ell_m-1}}, \dots, \mathbf{ct}_{m_{2,0}}] \in [\text{LWE}_s(m_{2,\ell_m-1}), \dots, \text{LWE}_s(m_{2,0})] \end{cases}$
 PUB : Public materials for **PBS-KS** and for **CBS**

Output: $\begin{cases} \mathbf{ct}_{f_{res}} = \begin{cases} \text{LWE}_s(s_{res}) \\ \mathbf{ct}_{e_{res}} = [ct_{e_{res}}, \dots, ct_{e_0}] \in [\text{LWE}_s(e_{res,\ell_e-1}), \dots, \text{LWE}_s(e_{res,0})] \\ \mathbf{ct}_{m_{res}} = [\text{LWE}_s(m_{res,\ell_m-1}), \dots, \text{LWE}_s(m_{res,0})] \end{cases} \end{cases}$

- 1 $(\mathbf{ct}_{m'_1}, \mathbf{ct}_{m'_2}, \mathbf{ct}_e) \leftarrow \text{AlignMantissa}(\mathbf{ct}_{e_1}, \mathbf{ct}_{m_1}, \mathbf{ct}_{e_2}, \mathbf{ct}_{m_2}, \text{PUB})$; /* Algo 7 */
- 2 $\mathbf{ct}_{m_{add}} \leftarrow \mathbf{ct}_{m'_1} + \mathbf{ct}_{m'_2}$
- 3 $(\mathbf{ct}_{m_{sub}}, \mathbf{ct}_{e_{sub}}, \mathbf{ct}_{s_{sub}}) \leftarrow \text{SubMantissa}(\mathbf{ct}_{m'_1}, \mathbf{ct}_{m'_2}, \mathbf{ct}_e, \mathbf{ct}_{s_1}, \text{PUB})$; /* Algo 8 */
- /* $\overline{\overline{\text{CT}_s}} \in \text{GGSW}_S^{\beta,\ell}(0)$ if $\mathbf{ct}_{s_1} + \mathbf{ct}_{s_2} \in \text{LWE}_s(0)$; $\overline{\overline{\text{CT}_s}} \in \text{GGSW}_S^{\beta,\ell}(1)$ otherwise. */
- 4 $\overline{\overline{\text{CT}_s}} = \text{CBS}(\mathbf{ct}_{s_1} + \mathbf{ct}_{s_2}, \text{PUB})$
- 5 $\mathbf{ct}_{e_{res}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_e, \mathbf{ct}_{e_{sub}}, \overline{\overline{\text{CT}_s}}),$
 $\mathbf{ct}_{m_{res}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{m_{add}}, \mathbf{ct}_{m_{sub}}, \overline{\overline{\text{CT}_s}});$ /* Algo 5 */
- 6 $\mathbf{ct}_{s_{res}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{s_1}, \mathbf{ct}_{s_{sub}}, \overline{\overline{\text{CT}_s}});$ /* Algo 5 */
- 7 **return** $\mathbf{ct}_{f_{res}} = (\mathbf{ct}_{s_{res}}, \mathbf{ct}_{e_{res}}, \mathbf{ct}_{m_{res}})$

Proof 11 (Correctness of Addition (Algorithm 9)) As defined in *AlignMantissa* (Algorithm 7), line 1 returns $\mathbf{e} = \max(\mathbf{e}_1, \mathbf{e}_2)$ and $\mathbf{ct}_{m'_1}$ and $\mathbf{ct}_{m'_2}$ aligned (such that if $\mathbf{e}_1 \geq \mathbf{e}_2$, $\mathbf{m}'_1 = \mathbf{m}_1$ and $\mathbf{m}'_2 = \lfloor \mathbf{m}_2 / 2^{\rho_m \cdot \gamma} \rfloor$ with $\gamma = \mathbf{e}_1 - \mathbf{e}_2$. And if $\mathbf{e}_1 \leq \mathbf{e}_2$, $\mathbf{m}'_2 = \mathbf{m}_2$ and $\mathbf{m}'_1 = \lfloor \mathbf{m}_1 / 2^{\rho_m \cdot \gamma} \rfloor$ with $\gamma = \mathbf{e}_2 - \mathbf{e}_1$).

Line 2 adds the two aligned mantissas \mathbf{m}'_1 and \mathbf{m}'_2 by adding together each LWE ciphertext. As the carry block is empty, these operations can be done directly.

As defined in Algorithm 8, line 3 returns $\mathbf{ct}_{m_{sub}} = |\mathbf{m}'_1 - \mathbf{m}'_2| \cdot 2^{\ell_m \cdot \alpha}$, $\mathbf{s}_{sub} = \mathbf{s}_1$ if $\mathbf{m}'_1 > \mathbf{m}'_2$, or $\mathbf{s}_{sub} = \mathbf{s}_2$ if $\mathbf{m}'_1 < \mathbf{m}'_2$. Assuming $\mathbf{m}'_1 \neq \mathbf{m}'_2$, let α be the index of the first non zero block of $\mathbf{m}_{res} = |\mathbf{m}'_1 - \mathbf{m}'_2| \cdot 2^{\ell_m \cdot \alpha}$, then if $\mathbf{e} > \alpha$ the algorithm returns $\mathbf{e}_{sub} = \mathbf{e} - \alpha$. If $\mathbf{m}'_1 = \mathbf{m}'_2$ or $\mathbf{e} < \alpha$, $\mathbf{m}_{sub} = 0$, $\mathbf{e}_{sub} = 0$ and $\mathbf{s}_{sub} = \mathbf{s}_1$.




Figure 4: This figure gives an overview of each step needed to perform the homomorphic floating-point addition.

By adding the sign on line 4 and performing a CBS over the result, we obtain a GGSW ciphertext which encrypt $\mathbf{s}_1 + \mathbf{s}_2$. So $\overline{\overline{\text{CT}}} \in \text{GGSW}_{\mathbf{S}}^{\beta, \ell}(0)$ if the two signs are equal or $\overline{\overline{\text{CT}}} \in \text{GGSW}_{\mathbf{S}}^{\beta, \ell}(1)$ if the signs are different.

So with this GGSW ciphertext, we will return the mantissa, the exponent and the sign corresponding to the addition if the signs are equal and if they are different, we return the mantissa, the exponent and the sign corresponding to the subtraction as proposed on the lemma 11.

Lemma 12 (Noise Constraints of Algorithm 9) *The output noise variances of ciphertexts of Algorithm 9, $\mathbf{ct}_{\mathbf{m}_{\text{res}}}$, $\mathbf{ct}_{\mathbf{e}_{\text{res}}}$ and $\mathbf{ct}_{\mathbf{s}_{\text{res}}}$, are respectively $\max(2\sigma_{\text{BR}}^2 + (\rho_e \ell_e + 5) \cdot \sigma_{\text{cmux}}^2, \sigma_{\text{BR}}^2 + (\ell_m + 1) \cdot \sigma_{\text{cmux}}^2) + \sigma_{\text{cmux}}^2, \sigma_{\text{BR}}^2 + 2\sigma_{\text{cmux}}^2$ and $5\sigma_{\text{BR}}^2 + \sigma_{\text{cmux}}^2$.*

To guarantee correctness of this operation, we need to find parameters that verify the following inequalities:

$$2 \max((\ell_m - 1) \cdot \sigma_{\text{cmux}}^2, \sigma_{\text{BR}}^2 + 2 \cdot \sigma_{\text{cmux}}^2) + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$$

$$\max(2\sigma_{\text{BR}}^2 + (\rho_e \ell_e + 5) \sigma_{\text{cmux}}^2, \sigma_{\text{BR}}^2 + (\ell_m + 1) \sigma_{\text{cmux}}^2) + \sigma_{\text{cmux}}^2 + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2.$$

With $\sigma_{\text{in},e}$ the noise variance of the input exponent ciphertexts, $\sigma_{\text{in},m}$ the noise variance of the input mantissa ciphertexts and with σ_{BR} the noise added by the blind rotation, σ_{KS} the noise added by the keyswitch, σ_{cmux} the noise added by the CMux and finally σ_{MS} , the noise added by the modulus switch.

Proof 12 (Proof of Lemma 12) Let us assume that the inputs of this algorithm are the outputs of Algorithm 6. It means that the variances of \mathbf{ct}_{s_i} , \mathbf{ct}_{e_i} and \mathbf{ct}_{m_i} are respectively σ_{BR}^2 , $\sigma_{\text{BR}}^2 + \sigma_{\text{cmux}}^2$ and $\sigma_{\text{BR}}^2 + \sigma_{\text{cmux}}^2$ (see Lemma 6).

The first line calls Algorithm 7, we use Lemma 8 to estimate the noise variances of $\mathbf{ct}_{m'_1}$, $\mathbf{ct}_{m'_2}$ which are equal to $\sigma_{\text{BR}}^2 + \left(\frac{\rho_e \ell_e + 5}{2}\right) \cdot \sigma_{\text{cmux}}^2$ and the noise variance of $\mathbf{ct}_{e'}$ which is equal to $\sigma_{\text{BR}}^2 + 2 \cdot \sigma_{\text{cmux}}^2$. Then, $\mathbf{ct}_{m'_1}$ and $\mathbf{ct}_{m'_2}$ are added which doubles the variance.

Next, we call Algorithm 8 and using Lemma 10, we deduce that the noise variances of $\mathbf{ct}_{m_{\text{sub}}}$, $\mathbf{ct}_{e_{\text{sub}}}$ and $\mathbf{ct}_{s_{\text{sub}}}$ are respectively $\sigma_{\text{BR}}^2 + (\ell_m + 1) \cdot \sigma_{\text{cmux}}^2$, $\sigma_{\text{BR}}^2 + \sigma_{\text{cmux}}^2$ and $5\sigma_{\text{BR}}^2$.

Then, we have a circuit bootstrap which must satisfies the following constraint $2\sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2$. Finally, we have an extended cmux for the mantissa, the exponent and the sign. The noise variances of $\mathbf{ct}_{m_{\text{res}}}$, $\mathbf{ct}_{e_{\text{res}}}$ and $\mathbf{ct}_{s_{\text{res}}}$ are respectively $\max(2\sigma_{\text{BR}}^2 + (\rho_e \ell_e + 5) \cdot \sigma_{\text{cmux}}^2, \sigma_{\text{BR}}^2 + (\ell_m + 1) \cdot \sigma_{\text{cmux}}^2) + \sigma_{\text{cmux}}^2$, $\sigma_{\text{BR}}^2 + 2\sigma_{\text{cmux}}^2$ and $5\sigma_{\text{BR}}^2 + \sigma_{\text{cmux}}^2$.

Using Lemmas 6, 2, 8 and 10 and by noticing that some of the inequalities are dominated by others, we find the complete set of constraints. As we want to be able to chain several additions, we will assume that after the addition, we apply Algorithm 6. The non-dominated set of constraints is the following:

$$\begin{aligned} & 2 \max((\ell_m - 1) \cdot \sigma_{\text{cmux}}^2, \sigma_{\text{BR}}^2 + 2 \cdot \sigma_{\text{cmux}}^2) + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2 \\ & \max(2\sigma_{\text{BR}}^2 + (\rho_e \ell_e + 5) \sigma_{\text{cmux}}^2, \sigma_{\text{BR}}^2 + (\ell_m + 1) \sigma_{\text{cmux}}^2) + \sigma_{\text{cmux}}^2 + \sigma_{\text{BR}}^2 + \sigma_{\text{KS}}^2 + \sigma_{\text{MS}}^2 \leq t^2. \end{aligned}$$

We need to find parameters that verify these inequalities to guarantee correctness.

5 Multiplication and Division

In this section, we introduce a very efficient floating-point multiplication with the HFP representation. Then, we detail a second algorithm to perform division. Finally, we briefly describe how to perform the ReLU and an approximated Sigmoid.

5.1 Multiplication

This operation computes the product of two HFPs. Following this procedure, it is necessary to apply the CarryPropagateFloat algorithm (Algorithm 6). This step ensures that the homomorphic floating-point number is returned to its standard representation, aligning it with the conventional HFP formalism. At a high level, the goal is to multiply the two mantissas without losing precision. Next, the exponent is updated by computing the sum of the two exponents and subtracting the bias. We note that we selected this bias to allow for efficient computation of this step. Due to the representation, we know that if the most significant ciphertext of the mantissa equals zero, or if the exponent is negative, the result of the operation is zero since we do not work with subnormal values. Otherwise, we return the result of the multiplication along with the updated exponent.

Lemma 13 (Multiplication (Algorithm 10)) *Let \mathbf{ct}_{f_i} such that $\mathbf{ct}_{s_i} \in \text{LWE}_s(s_i)$, $\mathbf{ct}_{e_i} = [\mathbf{ct}_{e_{i,\ell_e-1}}, \dots, \mathbf{ct}_{e_{i,0}}] \in [\text{LWE}_s(e_{i,\ell_e-1}), \dots, \text{LWE}_s(e_{i,0})]$ encrypting $e_i < 2^{\rho_e}$ and $\mathbf{ct}_{m_i} = [\mathbf{ct}_{m_{i,\ell_m-1}}, \dots, \mathbf{ct}_{m_{i,0}}] \in [\text{LWE}_s(m_{i,\ell_m-1}), \dots, \text{LWE}_s(m_{i,0})]$ encrypting $m_i < 2^{\rho_m}$ with $i \in \{1, 2\}$ be two ciphertexts encrypting $f_i = (-1)^{s_i} \cdot m_i \cdot (2^{\rho_m})^{e_i - \text{bias}}$. Let $(\mathbf{ct}_{m_{\text{res}}}, \mathbf{ct}_{e_{\text{res}}}, \mathbf{ct}_{s_{\text{res}}}) = \mathbf{ct}_{f_{\text{res}}} \leftarrow \text{Multiplication}(\mathbf{ct}_{f_1}, \mathbf{ct}_{f_2}, \text{PUB})$.*

Then $\text{DecryptFloat}(\mathbf{ct}_{f_{\text{res}}}) = f_{\text{res}} = (-1)^{s_{\text{res}}} \cdot m_{\text{res}} \cdot (2^{\rho_m})^{e_{\text{res}} - \text{bias}} = f_1 \cdot f_2 + \epsilon$ such that $\mathbf{ct}_{s_{\text{res}}} = \mathbf{ct}_{s_1} + \mathbf{ct}_{s_2}$ and ϵ is the maximum error added by the operation as express in Definition 5. If $m_1 \neq 0$ and $m_2 \neq 0$, if $e_1 + e_2 \geq \text{bias} - \ell_m + 1$ then $m_{\text{res}} = \lfloor m_1 \cdot m_2 / 2^{(\ell_m-1) \cdot \rho_m} \rfloor$ and $e_{\text{res}} = e_1 + e_2 - \text{bias} + \ell_m - 1$ with $|\epsilon| < (2^{\rho_m})^{e_{\text{res}} - \text{bias}}$. If $e_1 + e_2 < \text{bias} - \ell_m + 1$ then $m_{\text{res}} = 0$ and $e_{\text{res}} = 0$ with $|\epsilon| < 2^{\rho_m \cdot (\ell_m-1)} (2^{\rho_m})^{-\text{bias}}$. If $m_1 = 0$ or $m_2 = 0$ then $m_{\text{res}} = 0$ and $e_{\text{res}} = 0$ with $|\epsilon| < 2^{\rho_m \cdot (\ell_m-1)} (2^{\rho_m})^{-\text{bias}}$.

The complexity of the algorithm is: $\mathbb{C}_{\text{Multiplication}} = \mathbb{C}_{\text{IntMul}}^{\ell_m} + 2 \cdot \mathbb{C}_{\text{PBS}} + \mathbb{C}_{\text{CBS}}$.

Lemma 14 ($\mathbb{C}_{\text{IntMul}}^{\ell_m}$) *To simplify the algorithm, the operation IntMul proposed in the algorithm, have a complexity which can be bound by $(2 \cdot \ell_m^2 + \ell_m^2 / \rho_m) \cdot \mathbb{C}_{\text{PBS}}$. However, in our implementation, we use a slight modification of this algorithm to remove unnecessary computations (the part of the multiplication which does not appear in the final mantissa). In practice, the complexity is bounded by $(2 \cdot (\ell_m/2 + 1)^2 + \ell_m \cdot (\ell_m/2 + 1) / \rho_m) \cdot \mathbb{C}_{\text{PBS}}$.*

Proof 13 (Proof of complexity $\mathbb{C}_{\text{IntMul}}^{\ell_m}$) *The mantissa m represents a value in $[2^{\rho_m \cdot (\ell_m-1)}, 2^{\rho_m \cdot \ell_m} - 1]$. So m^2 is in*

$[2^{2 \cdot \rho_m \cdot (\ell_m - 1)}, 2^{2 \rho_m \cdot \ell_m} - 2^{\rho_m \cdot \ell_m + 1} + 1]$. After a multiplication, the smallest reachable value is $2^{2 \cdot \rho_m \cdot (\ell_m - 1)}$ and as we keep only the ℓ_m most significant blocks, all the values smaller than $2^{2 \cdot \rho_m \cdot (\ell_m - 1) - \ell_m \cdot \rho_m} = 2^{\rho_m \cdot \ell_m - 2 \rho_m}$ are lost so we don't need to compute them. These values correspond to the part of the mantissa \mathbf{m}' such that $\mathbf{m}'^2 < 2^{\rho_m \cdot \ell_m - 2 \rho_m}$ so the part \mathbf{m}' such that $\mathbf{m}' < 2^{\rho_m \cdot (\ell_m - 2)/2}$. The part of the mantissa \mathbf{m}' corresponds to the $\ell_m/2 - 1$ less significant blocks of the mantissa \mathbf{m} . So we only need to compute the multiplication on the $\ell_m/2 + 1$ most significant blocks.

Proof 14 (Correctness of Multiplication (Algorithm 10)) In the first step of the operation, we compute the sign of the multiplication by summing the two signs \mathbf{s}_1 and \mathbf{s}_2 . As the sign is on the padding bit, the addition is done modulo 2. If the signs are equal, after the operation, the sign is positive ($\mathbf{ct}_{\mathbf{s}_1} + \mathbf{ct}_{\mathbf{s}_2} \in \text{LWE}_s(0)$), otherwise the sign is negative ($\mathbf{ct}_{\mathbf{s}_1} + \mathbf{ct}_{\mathbf{s}_2} \in \text{LWE}_s(1)$).

Next, we compute the multiplication of the two mantissas. `IntMul` return the product of two integers. If $\mathbf{m}_1 \neq 0$ and $\mathbf{m}_2 \neq 0$, we have $\mathbf{m}_i \in [2^{(\ell_m - 1) \cdot \rho_m}, 2^{\ell_m \cdot \rho_m} - 1]$ for $i \in \{0, 1\}$, so $\mathbf{m}_1 \cdot \mathbf{m}_2 \in [2^{2 \cdot (\ell_m - 1) \cdot \rho_m}, 2^{2 \cdot \ell_m \cdot \rho_m} - 2^{\ell_m \cdot \rho_m + 1} + 1]$ (note that the value is stored in $2 \cdot \ell_m$ blocks). As we want to keep the classical representation of the mantissa (ℓ_m blocks, where the most significant block is non-zero except when the result equals zero), we will remove the least significant blocks, ensuring that only the ℓ_m most significant blocks remain.

To do so, we distinguish two cases after the multiplication. The case where $\mathbf{m}_1 \cdot \mathbf{m}_2 \in [2^{2 \cdot (\ell_m - 1) \cdot \rho_m}, 2^{(2 \cdot \ell_m - 1) \cdot \rho_m})$ (the most significant block after the multiplication contains zero) and the case where $\mathbf{m}_1 \cdot \mathbf{m}_2 \in [2^{(2 \cdot \ell_m - 1) \cdot \rho_m}, 2^{2 \cdot \ell_m \cdot \rho_m} - 2^{\ell_m \cdot \rho_m + 1} + 1]$. During the operation `IntMul`, the carry buffer of each ciphertext in $\mathbf{ct}_{\mathbf{m}_{\text{mul}}}$ is emptied. As the carry buffer of $\mathbf{ct}_{\mathbf{m}_{\text{mul}}, 2\ell_m - 1}$ and $\mathbf{ct}_{\mathbf{m}_{\text{mul}}, 2\ell_m - 2}$ are empty, by multiplying $\mathbf{ct}_{\mathbf{m}_{\text{mul}}, 2\ell_m - 1}$ by 2^{ρ_m} , we have $\mathbf{ct}_{\mathbf{m}_{\text{mul}}, 2\ell_m - 1} \cdot 2^{\rho_m} \in \{0\} \cup [2^{\rho_m}, 2^{2 \cdot \rho_m})$. As $\mathbf{ct}_{\mathbf{m}_{\text{mul}}, 2\ell_m - 2} \in [0, 2^{\rho_m})$, we can sum these two values such that $\mathbf{ct}_{\mathbf{m}_{\text{mul}}, 2\ell_m - 1} = \mathbf{ct}_{\mathbf{m}_{\text{mul}}, 2\ell_m - 1} \cdot 2^{\rho_m} + \mathbf{ct}_{\mathbf{m}_{\text{mul}}, 2\ell_m - 2}$ (Line 3). We have now $\mathbf{m}_1 \cdot \mathbf{m}_2 \in [2^{2 \cdot (\ell_m - 1) \cdot \rho_m}, 2^{2 \cdot \ell_m \cdot \rho_m} - 2^{\ell_m \cdot \rho_m + 1} + 1]$ stored in $2 \cdot \ell_m - 1$ blocks. Now, we remove the $\ell_m - 1$ less significant block of the multiplication which represents too small values for the mantissa precision and the most significant block which have its information already stored in the second most significant block $\mathbf{ct}_{\mathbf{m}_{\text{mul}}, 2\ell_m - 1}$. We obtain $\mathbf{m}_{\text{res}} = \lfloor \mathbf{m}_1 \cdot \mathbf{m}_2 / 2^{(\ell_m - 1) \cdot \rho_m} \rfloor \in [2^{(\ell_m - 1) \cdot \rho_m}, 2^{(\ell_m + 1) \cdot \rho_m})$ (Line 4). (In practice, we have modified the algorithm

Algorithm 10: $\text{CT}_{\text{fres}} \leftarrow \text{Multiplication}(\text{ct}_{\text{f1}}, \text{ct}_{\text{f2}}, \text{PUB})$

Context: $\begin{cases} \text{LUT}_m: & \text{LUT to return 1 if the value equals 0; 0 otherwise} \\ \text{LUT}_e: & \text{LUT to return 0 if } (x \geq 2^{\rho_e-1}); 1 \text{ otherwise} \end{cases}$

Input: $\begin{cases} \text{ct}_{\text{f1}} = \begin{cases} \text{ct}_{\text{s1}} \in \text{LWE}_s(\text{s1}) \\ \text{ct}_{\text{e1}} = [\text{ct}_{\text{e1}, \ell_e-1}, \dots, \text{ct}_{\text{e1}, 0}] \in [\text{LWE}_s(\text{e}_{1, \ell_e-1}), \dots, \text{LWE}_s(\text{e}_{1, 0})] \\ \text{ct}_{\text{m1}} = [\text{ct}_{\text{m1}, \ell_m-1}, \dots, \text{ct}_{\text{m1}, 0}] \in [\text{LWE}_s(\text{m}_{1, \ell_m-1}), \dots, \text{LWE}_s(\text{m}_{1, 0})] \end{cases} \\ \text{ct}_{\text{f2}} = \begin{cases} \text{ct}_{\text{s2}} \in \text{LWE}_s(\text{s2}) \\ \text{ct}_{\text{e2}} = [\text{ct}_{\text{e2}, \ell_e-1}, \dots, \text{ct}_{\text{e2}, 0}] \in [\text{LWE}_s(\text{e}_{2, \ell_e-1}), \dots, \text{LWE}_s(\text{e}_{2, 0})] \\ \text{ct}_{\text{m2}} = [\text{ct}_{\text{m2}, \ell_m-1}, \dots, \text{ct}_{\text{m2}, 0}] \in [\text{LWE}_s(\text{m}_{2, \ell_m-1}), \dots, \text{LWE}_s(\text{m}_{2, 0})] \end{cases} \end{cases}$

Output: $\begin{cases} \text{ct}_{\text{fres}} = \begin{cases} \text{ct}_{\text{sres}} \in \text{LWE}_s(\text{sres}) \\ \text{ct}_{\text{eres}} = [\text{ct}_{\text{eres}, \ell_e-1}, \dots, \text{ct}_{\text{eres}, 0}] \in [\text{LWE}_s(\text{e}_{\text{res}, \ell_e-1}), \dots, \text{LWE}_s(\text{e}_{\text{res}, 0})] \\ \text{ct}_{\text{mres}} = [\text{ct}_{\text{mres}, \ell_m-1}, \dots, \text{ct}_{\text{mres}, 0}] \in [\text{LWE}_s(\text{m}_{\text{res}, \ell_m-1}), \dots, \text{LWE}_s(\text{m}_{\text{res}, 0})] \end{cases} \end{cases}$

PUB : Public materials for **PBS-KS** and for **CBS**

- 1 $\text{ct}_{\text{sres}} \leftarrow \text{ct}_{\text{s1}} + \text{ct}_{\text{s2}}$
- 2 $\text{ct}_{\text{m}_{\text{mul}}} = [\text{ct}_{\text{m}_{\text{mul}}, 2\ell_m-1}, \dots, \text{ct}_{\text{m}_{\text{mul}}, 0}] \leftarrow \text{IntMul}(\text{ct}_{\text{m1}}, \text{ct}_{\text{m2}}, \text{PUB})$
- 3 $\text{ct}_{\text{m}_{\text{mul}}, 2\ell_m-2} \leftarrow \text{ct}_{\text{m}_{\text{mul}}, 2\ell_m-1} \cdot 2^{\rho_m} + \text{ct}_{\text{m}_{\text{mul}}, 2\ell_m-2}$
- 4 $\text{ct}_{\text{m}_{\text{res}}} = [\text{ct}_{\text{m}_{\text{res}}, \ell_m-1}, \dots, \text{ct}_{\text{m}_{\text{res}}, 0}] \leftarrow [\text{ct}_{\text{m}_{\text{mul}}, 2\ell_m-2}, \dots, \text{ct}_{\text{m}_{\text{mul}}, \ell_m-1}]$
 /* $\text{ct}_{\text{tmp}_m} \in \text{LWE}_s(1)$ if $\text{ct}_{\text{m}_{\text{res}}, \ell_m-1} \in \text{LWE}_s(0)$; $\text{ct}_{\text{tmp}_e} \in \text{LWE}_s(0)$ otherwise */
- 5 $\text{ct}_{\text{tmp}_m} \leftarrow \text{KS-PBS}(\text{ct}_{\text{m}_{\text{res}}, \ell_m-1}, \text{LUT}_m, \text{PUB})$
- 6 $\text{ct}_{\text{eres}} = [\text{ct}_{\text{eres}, \ell_e-1}, \dots, \text{ct}_{\text{eres}, 0}] \leftarrow \text{IntAdd}(\text{ct}_{\text{e1}}, \text{ct}_{\text{e2}}, \text{PUB})$
 /* $\text{ct}_{\text{tmp}_e} \in \text{LWE}_s(0)$ if $\text{ct}_{\text{eres}, \ell_e-1} \in \text{LWE}_s(x)$ with $x \geq 2^{\rho_e-1}$; $\text{ct}_{\text{tmp}_e} \in \text{LWE}_s(1)$ otherwise */
- 7 $\text{ct}_{\text{tmp}_e} \leftarrow \text{KS-PBS}(\text{ct}_{\text{eres}, \ell_e-1}, \text{LUT}_e, \text{PUB})$
 /* bias have been choose such that bias $-\ell_m+1$ is equal to $2^{\ell_e \cdot \rho_e-1}$ so this subtraction can be done only on the most significant block */
- 8 $\text{ct}_{\text{eres}, \ell_e-1} \leftarrow \text{ct}_{\text{eres}, \ell_e-1} - \text{TrivialEncrypt}(2^{\rho_e-1}, 1)$
 /* $\text{ct}_{\text{tmp}} \in \text{LWE}_s(0)$ if e is big enough and the most significant mantissa LWE is not null */
- 9 $\text{ct}_{\text{tmp}} \leftarrow \text{ct}_{\text{tmp}_m} + \text{ct}_{\text{tmp}_e}$
 /* encrypt 0 if $\text{ct}_{\text{tmp}} \in \text{LWE}_s(0)$, 1 otherwise */
- 10 $\overline{\text{CT}} \leftarrow \text{CBS}(\text{ct}_{\text{tmp}}, \text{PUB})$
- 11 $\text{ct}_{\text{eres}} \leftarrow \text{ExtendedCMux}(\text{ct}_{\text{eres}}, \text{TrivialEncrypt}(0, \ell_e), \overline{\text{CT}}),$
 $\text{ct}_{\text{m}_{\text{res}}} \leftarrow \text{ExtendedCMux}(\text{ct}_{\text{m}_{\text{res}}}, \text{TrivialEncrypt}(0, \ell_m), \overline{\text{CT}});$ /* Algo 5 */
- 12 **return** $\text{ct}_{\text{res}} = (\text{ct}_{\text{sres}}, \text{ct}_{\text{eres}}, \text{CT}_{\text{mres}})$

IntMul such that the carry propagation of the most significant block is not done and such that the useless parts of the multiplication are not computed

(Proof 13)). In the special case where $\mathbf{m}_1 = 0$ or $\mathbf{m}_2 = 0$, the previous step has no impact and `IntMul` will return the value 0 on each block. To distinguish the two cases, line 5, a PBS checks if the most significant block of the mantissa is equal to zero, such that $\mathbf{ct}_{\text{tmp}_m}$ is in $\text{LWE}_s(1)$ if $\mathbf{ct}_{\mathbf{m}_{\text{res}}, \ell_m - 1}$ is in $\text{LWE}_s(0)$, otherwise it returns $\mathbf{ct}_{\text{tmp}_m}$ in $\text{LWE}_s(0)$.

Now, we need to update the exponent \mathbf{e} . First we will sum the two exponents (Line 6). Next, as the exponent has the shape $\mathbf{e}'_i + \text{bias}$ after the sum we obtain $\mathbf{e}_1 + \mathbf{e}_2 = \mathbf{e}_{\text{res}} = \mathbf{e}'_{\text{res}} + 2 \cdot \text{bias}$. So to keep the same representation, we need to remove one bias. Moreover, in the previous step, we have removed the $\ell_m - 1$ blocks of the mantissa, so we need to add $\ell_m - 1$ to the exponent. In Section 2.5, we chose the bias such that $\text{bias} - \ell_m + 1 = 2^{\ell_e \cdot \rho_e - 1}$ so we only need to subtract $2^{\ell_e \cdot \rho_e - 1}$ from the sum of the exponents. Line 4, a PBS checks if the exponent is big enough and return an LWE ciphertext such that $\mathbf{ct}_{\text{tmp}_e} \in \text{LWE}_s(0)$ if $\mathbf{ct}_{\mathbf{e}_{\text{res}}, \ell_e - 1}$ is in $\text{LWE}_s(x)$ with $x \geq 2^{\rho_e - 1}$, otherwise it returns $\mathbf{ct}_{\text{tmp}_e} \in \text{LWE}_s(1)$. We can now subtract $2^{\ell_e \cdot \rho_e - 1}$ from the sum of the exponent (Line 8). This operation impacts only the most significant block of the exponent and can be performed directly.

Now, looking at the two previous control LWE ciphertexts ($\mathbf{ct}_{\text{tmp}_m}$ and $\mathbf{ct}_{\text{tmp}_e}$), by summing these two values, we obtain $\mathbf{ct}_{\text{tmp}_m} + \mathbf{ct}_{\text{tmp}_e} \in \text{LWE}_s(0)$ if $\mathbf{ct}_{\mathbf{m}_{\text{res}}, \ell_m - 1} \notin \text{LWE}_s(0)$ and $\mathbf{ct}_{\mathbf{e}_{\text{res}}, \ell_e - 1} \in \text{LWE}_s(x)$ with $x \geq 2^{\rho_e - 1}$. Otherwise, one of the two conditions to perform the multiplication is unmet and the multiplication is not feasible. By using a circuit bootstrapping, we obtain a GGSW ciphertext $\overline{\mathbf{CT}}$ such that $\overline{\mathbf{CT}} \in \text{GGSW}_S^{\beta, \ell}(0)$ if we can perform the multiplication and $\overline{\mathbf{CT}} \in \text{GGSW}_S^{\beta, \ell}(1)$ otherwise. With the last 2 lines, if $\overline{\mathbf{CT}}$ is in $\text{GGSW}_S^{\beta, \ell}(0)$, the algorithm returns the result of the multiplication, otherwise the multiplication is not doable and it returns zero.

5.2 Division

This operation computes the division of two HFPs. Following this procedure, it is necessary to apply the CarryPropagateFloat algorithm (Algorithm 6). This step will ensure that the homomorphic floating-point number is returned to its standard representation, aligning it with the conventional HFP formalism.

Lemma 15 (Division (Algorithm 11)) *Let $\mathbf{ct}_{\mathbf{f}_i}$ such that $\mathbf{ct}_{\mathbf{f}_i} \in \text{LWE}_s(\mathbf{f}_i)$, $\mathbf{ct}_{\mathbf{e}_i} = [\mathbf{ct}_{\mathbf{e}_i, \ell_e - 1}, \dots, \mathbf{ct}_{\mathbf{e}_i, 0}] \in [\text{LWE}_s(\mathbf{e}_i, \ell_e - 1), \dots, \text{LWE}_s(\mathbf{e}_i, 0)]$*

encrypting $\epsilon < 2^{\rho_\epsilon}$ and $\mathbf{ct}_{m_i} = [\mathbf{ct}_{m_i, \ell_m-1}, \dots, \mathbf{ct}_{m_i, 0}] \in [\text{LWE}_s(m_{i, \ell_m-1}), \dots, \text{LWE}_s(m_{i, 0})]$ encrypting $m < 2^{\rho_m}$ with $i \in \{1, 2\}$ be two ciphertexts encrypting $f_i = (-1)^{s_i} \cdot m_i \cdot (2^{\rho_m})^{\epsilon_i - \text{bias}}$. Let $(\mathbf{ct}_{m_{\text{res}}}, \mathbf{ct}_{\epsilon_{\text{res}}}, \mathbf{ct}_{s_{\text{res}}}) = \mathbf{ct}_{f_{\text{res}}} \leftarrow \text{Division}(\mathbf{ct}_{f_1}, \mathbf{ct}_{f_2}, \text{PUB})$. Then, $\text{DecryptFloat}(\mathbf{ct}_{f_{\text{res}}}) = f_{\text{res}} = (-1)^{s_{\text{res}}} \cdot m_{\text{res}} \cdot (2^{\rho_m})^{\epsilon_{\text{res}} - \text{bias}} = f_1/f_2 + \epsilon$ with $|\epsilon| < (2^{\rho_m})^{\epsilon_{\text{res}} - \text{bias}}$ s.t. $\mathbf{ct}_{s_{\text{res}}} = \mathbf{ct}_{s_1} + \mathbf{ct}_{s_2}$. If $\mathbf{ct}_{\epsilon_2} < \mathbf{ct}_{\epsilon_1} + \text{bias} + \ell_m - 1$, then $\mathbf{ct}_{m_{\text{res}}} = \lfloor \mathbf{ct}_{m_1} \cdot 2^{\ell_m-1} / \mathbf{ct}_{m_2} \rfloor$ and $\mathbf{ct}_{\epsilon_{\text{res}}} = \mathbf{ct}_{\epsilon_1} + \text{bias} + \ell_m - 1 - \mathbf{ct}_{\epsilon_2}$. Else, $\mathbf{ct}_{\epsilon_{\text{res}}} = 0$ and $\mathbf{ct}_{m_{\text{res}}} = 0$. The complexity of the algorithm is: $\mathbb{C}_{\text{Division}} = \mathbb{C}_{\text{IntSub}^*}^{\ell_\epsilon} + \mathbb{C}_{\text{IntDiv}}^{2 \cdot \ell_m} + \mathbb{C}_{\text{CBS}}$.

Algorithm 11: $\mathbf{ct}_f \leftarrow \text{Division}(\mathbf{ct}_{f_1}, \mathbf{ct}_{f_2}, \text{PUB})$

Input: $\begin{cases} \mathbf{ct}_{f_1} = \begin{cases} \mathbf{ct}_{s_1} \in \text{LWE}_s(s_1) \\ \mathbf{ct}_{\epsilon_1} = [\mathbf{ct}_{\epsilon_1, \ell_\epsilon-1}, \dots, \mathbf{ct}_{\epsilon_1, 0}] \in [\text{LWE}_s(\epsilon_{1, \ell_\epsilon-1}), \dots, \text{LWE}_s(\epsilon_{1, 0})] \\ \mathbf{ct}_{m_1} = [\mathbf{ct}_{m_1, \ell_m-1}, \dots, \mathbf{ct}_{m_1, 0}] \in [\text{LWE}_s(m_{1, \ell_m-1}), \dots, \text{LWE}_s(m_{1, 0})] \end{cases} \\ \mathbf{ct}_{f_2} = \begin{cases} \mathbf{ct}_{s_2} \in \text{LWE}_s(s_2) \\ \mathbf{ct}_{\epsilon_2} = [\mathbf{ct}_{\epsilon_2, \ell_\epsilon-1}, \dots, \mathbf{ct}_{\epsilon_2, 0}] \in [\text{LWE}_s(\epsilon_{2, \ell_\epsilon-1}), \dots, \text{LWE}_s(\epsilon_{2, 0})] \\ \mathbf{ct}_{m_2} = [\mathbf{ct}_{m_2, \ell_m-1}, \dots, \mathbf{ct}_{m_2, 0}] \in [\text{LWE}_s(m_{2, \ell_m-1}), \dots, \text{LWE}_s(m_{2, 0})] \end{cases} \\ \text{PUB} : \text{Public materials for PBS-KS and for CBS} \end{cases}$

Output: $\begin{cases} \mathbf{ct}_{f_{\text{res}}} = \begin{cases} \mathbf{ct}_{s_{\text{res}}} \in \text{LWE}_s(s_{\text{res}}) \\ \mathbf{ct}_{\epsilon_{\text{res}}} = [\mathbf{ct}_{\epsilon_{\text{res}}, \ell_\epsilon-1}, \dots, \mathbf{ct}_{\epsilon_{\text{res}}, 0}] \in [\text{LWE}_s(\epsilon_{\text{res}, \ell_\epsilon-1}), \dots, \text{LWE}_s(\epsilon_{\text{res}, 0})] \\ \mathbf{ct}_{m_{\text{res}}} = [\mathbf{ct}_{m_{\text{res}}, \ell_m-1}, \dots, \mathbf{ct}_{m_{\text{res}}, 0}] \in [\text{LWE}_s(m_{\text{res}, \ell_m-1}), \dots, \text{LWE}_s(m_{\text{res}, 0})] \end{cases} \end{cases}$

- 1 $\mathbf{ct}_{s_{\text{res}}} \leftarrow \mathbf{ct}_{s_1} + \mathbf{ct}_{s_2}$
- 2 $\mathbf{ct}_{\epsilon_1} \leftarrow \mathbf{ct}_{\epsilon_1} + \text{TrivialEncrypt}(\text{bias} + \ell_m - 1, \ell_\epsilon)$
- 3 $(\mathbf{ct}_{\epsilon_s}, \mathbf{ct}_\epsilon = [\mathbf{ct}_{\epsilon, \ell_\epsilon}, \dots, \mathbf{ct}_{\epsilon, 0}]) \leftarrow \text{IntSub}^*(\mathbf{ct}_{\epsilon_1}, \mathbf{ct}_{\epsilon_2}, \text{PUB})$
- 4 $\mathbf{ct}_{m_1} = [\mathbf{ct}_{m_1, \ell_m-1}, \dots, \mathbf{ct}_{m_1, 0}, \mathbf{ct}_0, \dots, \mathbf{ct}_0] \leftarrow \mathbf{ct}_{m_1} || \text{TrivialEncrypt}(0, \ell_m - 1)$
- 5 $\mathbf{ct}_{m_2} = [\mathbf{ct}_0, \dots, \mathbf{ct}_0, \mathbf{ct}_{m_2, \ell_m-1}, \dots, \mathbf{ct}_{m_2, 0}] \leftarrow \text{TrivialEncrypt}(0, \ell_m - 1) || \mathbf{ct}_{m_2}$
- 6 $\mathbf{ct}_{m_{\text{div}}} = [\mathbf{ct}_{m_{\text{div}}, 2 \cdot \ell_m-1}, \dots, \mathbf{ct}_{m_{\text{div}}, 0}] \leftarrow \text{IntDiv}(\mathbf{ct}_{m_1}, \mathbf{ct}_{m_2}, \text{PUB})$
- 7 $\mathbf{ct}_{m_{\text{div}}, 2 \ell_m-2} \leftarrow \mathbf{ct}_{m_{\text{div}}, 2 \ell_m-1} \cdot 2^{\rho_m} + \mathbf{ct}_{m_{\text{div}}, 2 \ell_m-2}$
- 8 $\mathbf{ct}_{m_{\text{res}}} = [\mathbf{ct}_{m_{\text{res}}, \ell_m-1}, \dots, \mathbf{ct}_{m_{\text{res}}, 0}] \leftarrow [\mathbf{ct}_{m_{\text{div}}, 2 \ell_m-2}, \dots, \mathbf{ct}_{m_{\text{div}}, \ell_m-1}]$
- 9 $\overline{\overline{\text{CT}}}_{\epsilon_s} \leftarrow \text{CBS}(\mathbf{ct}_{\epsilon_s}, \text{PUB})$
- 10 $\mathbf{ct}_{\epsilon_{\text{res}}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_\epsilon, \text{TrivialEncrypt}(0, \ell_\epsilon), \overline{\overline{\text{CT}}}_{\epsilon_s}),$
 $\mathbf{ct}_{m_{\text{res}}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{m_{\text{res}}}, \text{TrivialEncrypt}(0, \ell_m), \overline{\overline{\text{CT}}}_{\epsilon_s}); \quad /* \text{ Algo 5 } *$
 $*/$
- 11 **return** $\mathbf{ct}_f = (\mathbf{ct}_{s_{\text{res}}}, \mathbf{ct}_{\epsilon_{\text{res}}}, \mathbf{ct}_{m_{\text{res}}})$

Proof 15 (Correctness of Division (Algorithm 11)) *On the first line, we compute the output sign. If the signs are equal, the sign after the operation is positive ($\text{ct}_{s_1} + \text{ct}_{s_2} \in \text{LWE}_s(0)$), otherwise the sign is negative ($\text{ct}_{s_1} + \text{ct}_{s_2} \in \text{LWE}_s(1)$).*

The algorithm IntDiv takes two vectors of ciphertexts that represent two integers and returns the quotient of the division. In our context, we can not divide directly the mantissas. In fact, the mantissas are in the same interval and very close, if we were to divide them directly, the quotient would only be a value of a few digits. In our case, we want a value in the interval $[2^{\rho_m \cdot (\ell_m - 1)}, 2^{\rho_m \cdot \ell_m})$. To get a result in the right interval, we add some blocks encrypting zeros after the blocks of the first mantissa \mathbf{m}_1 (Line 4). Adding the zeros to $\text{ct}_{\mathbf{m}_1}$ corresponds to compute $\mathbf{m}_1 \cdot 2^{\ell_m \cdot \rho_m} \in [2^{\rho_m \cdot (2\ell_m - 1)}, 2^{\rho_m \cdot 2\ell_m})$. Now if we divide $\mathbf{m}_1 \cdot 2^{\ell_m \cdot \rho_m}$ by \mathbf{m}_2 , we will obtain a result in $[2^{(\ell_m - 1) \cdot \rho_m}, 2^{(\ell_m + 1) \cdot \rho_m})$. As explained in Proof 14, this value can be stored in ℓ_m blocks if we use the carry buffer of the most significant block. The carry buffer will be later cleaned during the call to CarryPropagateFloat (Algorithm 6). So after the shift of the mantissa (Line 4) and the division (Line 6), we obtain a new mantissa in the interval $[2^{(\ell_m - 1) \cdot \rho_m}, 2^{(\ell_m + 1) \cdot \rho_m})$.

After the division of the mantissa, we need to update the exponent. To do so, we need to subtract (Line 3) the two exponents, then add the bias and finally add the number of trivial ciphertexts added in Line 4. If the subtraction of the exponent returns a negative result (Line 3, $\text{ct}_e \in \text{LWE}_s(0)$), the division can not be done. In this case, Algorithm 11 returns the value zero (Line 10 and 11), otherwise it returns the result of the division of the two floating-point numbers.

6 Experimental Results

In this section, we demonstrate the practicability of our results by providing all cryptographic parameters, encodings, and both sequential and parallel timings.

Encodings In Table 3, we describe the different encodings used to represent 64, 32, 16 and 8 bits floating-point numbers (Sec. 3.2) in the homomorphic world. Research in [BBB⁺22] indicates that a 4-bit precision message leads to the best precision-cost ratio; therefore, we focus on representations with $\rho_m = \rho_e = 2$. However, variations with $\rho_m \neq \rho_e$ may yield better tim-

ings depending on the specific use case. Additionally, in Table 3, we give the encoding for the TFHE-minifloats encoded over 8 bits as detailed in Sec. 3.1. For the TFHE-Minifloats, the value of the **bias** does not impact the timings and can be freely chosen.

	ℓ_m	ρ_m	ℓ_c	ρ_c	bias
TFHE_FP _{64b}	27	2	5	2	539
TFHE_FP _{32b}	13	2	4	2	140
TFHE_FP _{16b}	6	2	3	2	37
TFHE_FP _{8b}	3	2	2	2	10
TFHE-Minifloat ^{$\rho=4$}	3	\emptyset	4	\emptyset	8
TFHE-Minifloat ^{$\rho=2$}	3	\emptyset	4	\emptyset	8

Table 3: Encodings for HFP and Minifloats

Parameter Selection In Lemma 12, we found two noise constraints that the parameters must satisfy in order to guarantee the correctness of Algorithm 9. We applied the same reasoning to Algorithm 10 and found that all the additional noise constraints are dominated by the constraints introduced in Lemma 12. It means that parameters that satisfy the constraints of Algorithm 9 will also satisfy the constraints of Algorithm 10.

As explained in Lemmas 11 and 13, the number of PBSs in each algorithm is different and this has an impact on the failure probability of each algorithm. We followed [BBB⁺22]’s blueprint to compute the individual PBS failure probability using the number of dominant PBS in each algorithm. Using the parameters presented in Table ??, the maximal failure probability for the homomorphic addition and for the homomorphic multiplication are respectively $2^{-13.9}$ and $2^{-12.8}$ (note that the failure probability of one KS-PBS is smaller than 2^{-40}). We tested these parameters on a chain of a hundred operations on random inputs with random operations without detecting any errors due to the noise, only errors due to floating-point approximations.

Timings All of our experiments have been carried out on AWS with a m6i.metal instance Intel Xeon 8375C (Ice Lake) at 3.5 GHz, with 128 vCPUs and 512.0 GiB of memory using the TFHE-rs library [Zam22]. Our code is available here¹. In Table 5, we give the timings in seconds for all the

¹https://github.com/zama-ai/tfhe-rs/tree/artifact_tches_2025

arithmetic operations (i.e., add, sub, mul, div) and the ReLU and Sigmoid functions. Both sequential and parallel timings are given when possible (e.g. the division over integers is only implemented in parallel in TFHE-rs). Note that all arithmetic operations are followed by a carry propagation, which is obviously taken into account in the timings.

Finally, in Table 6, we present the timings for the WoP-PBS based approach. Although the multiplication timings are quite similar between HFP and minifloats, the addition operations perform significantly better using the WoP-PBS. This means that when computing with 8-bit floats, using the minifloats is generally better. Any larger precision requires to run the HFP method, whose timings show that many circuits based on floats can be practically evaluated for the first time. Only the division operation cannot be considered as practical. Note that the division is suffering from the slowness of the division over the integers, and not really from the approach described in here.

		Addition (Alg.9 & 6)	Multiplication (Alg.10 & 6)	Division (Alg.11 & 6)	Sigmoid (Alg.13)	Relu (Alg.12)
TFHE_FP _{64b}	Sequential	12.32 s	87.15 s	∅	0.342 s	0.122 s
	Parallel	3.98 s	2.26 s	39.75 s	∅	∅
TFHE_FP _{32b}	Sequential	7.10 s	20.57 s	∅	0.342 s	0.120 s
	Parallel	2.50 s	1.03 s	15.18 s	∅	∅
TFHE_FP _{16b}	Sequential	3.89 s	3.83 s	∅	0.361 s	0.155 s
	Parallel	1.52 s	0.558 s	4.34 s	∅	∅
TFHE_FP _{8b}	Sequential	2.21 s	1.19 s	∅	0.388 s	0.153 s
	Parallel	1.13 s	0.444 s	1.76 s	∅	∅

Table 5: Timings of the HFP depending on the precision.

	TFHE-Minifloat ^{$\rho=4$}	TFHE-Minifloat ^{$\rho=2$}
Bivariate Operation (e.g., add, mul, ...)	1.2819 s	0.9957 s

Table 6: Timings for the 8 bit representations of the TFHE-Minifloat.

In the previous table, we present benchmarks obtained with a failure probability around 2^{-14} . To better evaluate our new algorithms, we also include benchmarks of the addition and the multiplication with a failure probability of around 2^{-40} (see Table 7). We observe that reducing the

failure probability has only a minor impact on execution time which proves that our contribution scales well with small failure probabilities.

		Add	Mul
TFHE_FP _{32b}	Sequential	7.49s	23.2s
TFHE_FP _{32b}	Parallelized	2.83s	1.24s
TFHE_FP _{64b}	Sequential	13.3s	98.2s
TFHE_FP _{64b}	Parallelized	4.28s	2.75s

Table 7: Performance for Addition and Multiplication using $\text{pfail} \leq 2^{-40}$

7 More Features over HFP

This section extends our approach to cover a wider range of practical applications by adding the support of special values and efficient approximate functions.

7.1 Managing Special Values

In the classical floating-point arithmetic, when a value overflows the highest bound of the exponent, this value can not be represented anymore. In this case, the floating-point reaches the infinity. As previously described, our algorithms do not manage the values plus/minus infinity and Not a Number (NaN), but as we show now, they can easily be extended to do so. The idea is to add two encrypted Booleans to represent $+\infty$ and $-\infty$ such that: if only one is set, then it means that we have reached plus (resp. minus) infinity; if both are set, the value is interpreted as NaN. During an operation, the infinity value is reached as soon as an overflow occurs on the exponent, i.e., if the carry of the most significant block of the exponent is not empty. This check is done by computing a simple PBS on $\text{ct}_{\text{res}, \ell_e - 1}$, which returns a flag encrypting 0 if the carry is empty, or 1 otherwise. This flag is then given as an additional input to the **CarryPropagateFloat** (Alg. 6). Then, by computing a **CBS** on the sign, this will return the correct sign of the flag, i.e., plus or minus infinity (or 0). This process ends with the computation of a simple CMux tree which will properly update Booleans ciphertext depending on the flag value. Regarding the support of special values, the overhead is linear in the number of blocks composing the HFP. Thus, in comparison with the

numbers of PBS or CBS needed to perform operations without any special values, the overhead should be negligible.

7.2 Computing Function Approximations

Beyond the arithmetic operations, floating-point numbers are particularly convenient to compute approximations of complex functions, via the Taylor series. A Taylor series of a real function $f(x)$ that is infinitely differentiable at a real number a is the power series $\sum_{n=0}^{\infty} \frac{f^{(n)}(a)}{n!} (x-a)^n$, where $n!$ denotes the factorial of n and $f^{(n)}(a)$ denotes the n -th derivative of f evaluated at the point a . When $a = 0$, this is called a Maclaurin serie and takes the form $\sum_{n=0}^{\infty} \frac{f^{(n)}(0)}{n!} x^n$. The Maclaurin method is advantageous when working with homomorphic floating-point numbers, given that the value of $f(a)$ is not known. Computing such a serie is a direct application of our method, as each of its term can be computed using the previously defined arithmetic operators. Another advantage is that the needed precision and the computational time can be adjusted to fit the use case, i.e., by changing the value of n . For instance, we have practically computed $\sin(x) \approx \sum_{n=0}^2 \frac{(-1)^n}{(2n+1)!} x^{2n+1} = x - \frac{x^3}{3!} + \frac{x^5}{5!}$ and $\cos(x) \approx \sum_{n=0}^2 \frac{(-1)^n}{(2n)!} x^{2n} = 1 - \frac{x^2}{2!} + \frac{x^4}{4!}$, which gives good results for values of $x \in [-1, 1]$. In Table 8, we present numerical values obtained from our approximations of the cosine and sine functions using the Maclaurin series.

	Cos(0.9636989235877991)	Sin(0.41880202293395996)
Exact value (64 bits)	0.5704859425112639	0.4066663011129846
Approximate value (64 bits)	0.5715802311897278	0.4066667483866177
Approximate value (32 bits)	0.57158023	0.40666676
This work (32bits)	0.5715802311897278	0.40666675567626953

Table 8: Result obtain for the cosinus and sinus with Maclaurin series. The bold digits are the one which are equal to the digits of the approximate result of the clear double precision value.

7.3 Other Operations

With the homomorphic floating-point representation, we can efficiently support usual functions used in machine learning. To evaluate the ReLU function, we can apply a circuit bootstrapping on the sign \mathfrak{s} and return either the

input floating-point or an encryption of zero using a cmux. The complete algorithm is detailed in Algorithm 12 in Section 7.3.

In the same manner, we can evaluate an approximate sigmoid function that returns the identity for values in the interval $[-1, 1]$, and returns the constant value 1 or -1 otherwise. More details are provided in Algorithm 13 in Section 7.3. To be closer to the classical sigmoid, we can combine this approximate sigmoid with the Maclaurin series introduced in Section 7.2. Other classical operations like the minimum, the maximum or the equality between two values can be easily performed on homomorphic floating-point numbers.

7.3.1 ReLU

Lemma 16 (ReLU (Algorithm 12)) *Let \mathbf{ct}_f such that $\mathbf{ct}_s \in \text{LWE}_s(s), \mathbf{ct}_e = [\mathbf{ct}_{e_{\ell_e-1}}, \dots, \mathbf{ct}_{e_0}] \in [\text{LWE}_s(e_{\ell_e-1}), \dots, \text{LWE}_s(e_0)]$ encrypting $e < 2^{\rho_e}$ and $\mathbf{ct}_m = [\mathbf{ct}_{m_{\ell_m-1}}, \dots, \mathbf{ct}_{m_0}] \in [\text{LWE}_s(m_{\ell_m-1}), \dots, \text{LWE}_s(m_0)]$ encrypting $m < 2^{\rho_m}$ be ciphertexts encrypting $f = (-1)^s \cdot m \cdot (2^{\rho_m})^{e-\text{bias}}$.*

Let $(\mathbf{ct}_{m_{\text{res}}}, \mathbf{ct}_{e_{\text{res}}}, \mathbf{ct}_{s_{\text{res}}}) = \mathbf{ct}_{f_{\text{res}}} \leftarrow \text{Relu}(\mathbf{ct}_f, \text{PUB})$. Then $\text{DecryptFloat}(\mathbf{ct}_{f_{\text{res}}}) = f_{\text{res}}$ with $f_{\text{res}} = (-1)^{s_{\text{res}}} \cdot m_{\text{res}} \cdot (2^{\rho_m})^{e_{\text{res}}-\text{bias}}$ such that if $\mathbf{ct}_s \in \text{LWE}_s(0)$, $\mathbf{ct}_{s_{\text{res}}} = \mathbf{ct}_s$, $\mathbf{ct}_{e_{\text{res}}} = \mathbf{ct}_e$ and $\mathbf{ct}_{m_{\text{res}}} = \mathbf{ct}_m$ Else $\mathbf{ct}_{f_{\text{res}}}$ encrypt zero.

The complexity of the algorithm is: $\mathbb{C}_{\text{Division}} = \mathbb{C}_{\text{IntSub}^}^{\ell_e} + \mathbb{C}_{\text{IntDiv}}^{2 \cdot \ell_m} + \mathbb{C}_{\text{CBS}}$*

Algorithm 12: $\mathbf{ct}_f \leftarrow \text{ReLU}(\mathbf{ct}_f, \text{PUB})$

Input: $\begin{cases} \mathbf{ct}_s \in \text{LWE}_s(s) \\ \mathbf{ct}_e = [\mathbf{ct}_{e_{\ell_e-1}}, \dots, \mathbf{ct}_{e_0}] \in [\text{LWE}_s(e_{\ell_e-1}), \dots, \text{LWE}_s(e_0)] \\ \mathbf{ct}_m = [\mathbf{ct}_{m_{\ell_m-1}}, \dots, \mathbf{ct}_{m_0}] \in [\text{LWE}_s(m_{\ell_m-1}), \dots, \text{LWE}_s(m_0)] \end{cases}$

PUB: Public materials for **PBS-KS** and for **CBS**

Output: $\begin{cases} \mathbf{ct}_s \in \text{LWE}_s(s_{\text{res}}) \\ \mathbf{ct}_{e_{\text{res}}} = [\mathbf{ct}_{e_{\text{res}, \ell_e-1}}, \dots, \mathbf{ct}_{e_{\text{res}, 0}}] \in [\text{LWE}_s(e_{\text{res}, \ell_e-1}), \dots, \text{LWE}_s(e_{\text{res}, 0})] \\ \mathbf{ct}_{m_{\text{res}}} = [\mathbf{ct}_{m_{\text{res}, \ell_m-1}}, \dots, \mathbf{ct}_{m_{\text{res}, 0}}] \in [\text{LWE}_s(m_{\text{res}, \ell_m-1}), \dots, \text{LWE}_s(m_{\text{res}, 0})] \end{cases}$

/ encrypt 0 if sign == 0, 1 otherwise */*

- 1 $\overline{\overline{\text{CT}}} \leftarrow \text{CBS}(\mathbf{ct}_s, \text{PUB})$
- 2 $\mathbf{ct}_{e_{\text{res}}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_e, \text{TrivialEncrypt}(0, \ell_e), \overline{\overline{\text{CT}}})$
- 3 $\mathbf{ct}_{m_{\text{res}}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_m, \text{TrivialEncrypt}(0, \ell_m), \overline{\overline{\text{CT}}})$
- 4 **return** $\mathbf{ct}_f = [\mathbf{ct}_s, \mathbf{ct}_{m_{\text{res}}}, \mathbf{ct}_{e_{\text{res}}}]$

Proof 16 (ReLU (Algorithm 12)) *First we use a CBS on the sign to obtain a GGSW ciphertext such that $\overline{\overline{\text{CT}}} \in \text{GGSW}_{\mathbf{s}}^{\beta, \ell}(0)$ if $\text{ct}_{\mathbf{s}}$ is in $\text{LWE}_{\mathbf{s}}(0)$, otherwise, $\overline{\overline{\text{CT}}}$ is in $\text{GGSW}_{\mathbf{s}}^{\beta, \ell}(1)$*

Next with the GGSW ciphertext, we return the input ciphertext if the sign is positive otherwise we return zero.

7.3.2 Approximate Sigmoid

An efficient algorithm to compute an approximation of the sigmoid function compatible with the HFP representation is presented in Algorithm 13.

Algorithm 13: $\mathbf{ct}_f \leftarrow \text{ApproxSigmoid}(\mathbf{ct}_f, \text{PUB})$

Context: $\begin{cases} \text{LUT}_m : & \text{LUT to return 1 if the value equals 1, 1 otherwise} \\ \text{LUT}_e : & \text{LUT to return 0 if } (x \geq 2^{\rho_e-1}); 1 \text{ otherwise} \end{cases}$

Input: $\begin{cases} \mathbf{ct}_s \in \text{LWE}_s(\mathfrak{s}) \\ \mathbf{ct}_e = [\mathbf{ct}_{e,\ell_e-1}, \dots, \mathbf{ct}_{e,0}] \in [\text{LWE}_s(\mathfrak{e}_{\ell_e-1}), \dots, \text{LWE}_s(\mathfrak{e}_0)] \\ \mathbf{ct}_m = [\mathbf{ct}_{m,\ell_m-1}, \dots, \mathbf{ct}_{m,0}] \in [\text{LWE}_s(\mathfrak{m}_{\ell_m-1}), \dots, \text{LWE}_s(\mathfrak{m}_0)] \end{cases}$

Output: $\begin{cases} \mathbf{ct}_f = \begin{cases} \mathbf{ct}_s \in \text{LWE}_s(\mathfrak{s}) \\ \mathbf{ct}_e = [\mathbf{ct}_{e,\ell_e-1}, \dots, \mathbf{ct}_{e,0}] \in [\text{LWE}_s(\mathfrak{e}_{\ell_e-1}), \dots, \text{LWE}_s(\mathfrak{e}_0)] \\ \mathbf{ct}_m = [\mathbf{ct}_{m,\ell_m-1}, \dots, \mathbf{ct}_{m,0}] \in [\text{LWE}_s(\mathfrak{m}_{\ell_m-1}), \dots, \text{LWE}_s(\mathfrak{m}_0)] \end{cases} \end{cases}$

PUB : Public materials for **PBS-KS** and forCBS

- 1 $\mathbf{ct}_1 = [\mathbf{ct}_{e_1}, \mathbf{ct}_{m_1}, \mathbf{ct}_{s_1}] \leftarrow \text{TrivialEncryptFloat}(1)$
- 2 $\mathbf{ct}_{-1} = [\mathbf{ct}_{e_{-1}}, \mathbf{ct}_{m_{-1}}, \mathbf{ct}_{s_{-1}}] \leftarrow \text{TrivialEncryptFloat}(-1)$
- 3 $\overline{\overline{\mathbf{CT}}}_s \leftarrow \text{CBS}(\mathbf{ct}_s, \text{PUB}); \quad /* \text{encrypt 0 if } sign == 0, 1 \text{ otherwise} */$
- 4 $\mathbf{ct}_{s_{tmp}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{s_1}, \mathbf{ct}_{s_{-1}}, \overline{\overline{\mathbf{CT}}}_s),$
 $\mathbf{ct}_{e_{tmp}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{e_1}, \mathbf{ct}_{e_{-1}}, \overline{\overline{\mathbf{CT}}}_s),$
 $\mathbf{ct}_{m_{tmp}} \leftarrow \text{ExtendedCMux}(\mathbf{ct}_{m_1}, \mathbf{ct}_{m_{-1}}, \overline{\overline{\mathbf{CT}}}_s)$
- 5 $\mathbf{ct}_{tmp_e} \leftarrow \text{KS-PBS}(\mathbf{ct}_{e_{tmp}}, \text{PUB}, \text{LUT}_e)$
 $/* \text{If } \mathbf{ct}_{e_{tmp}} \in \text{LWE}_s(x) \text{ with } x < 2^{\rho_e-1} \text{ then } \text{LWE}_s(1) \text{ else } \text{LWE}_s(0) */$
- 6 $\mathbf{ct}_{tmp_m} \leftarrow \text{KS-PBS}(\mathbf{ct}_{m_{tmp}}, \text{PUB}, \text{LUT}_m)$
 $/* \text{If } \mathbf{ct}_{m_{tmp}} \in \text{LWE}_s(x) \text{ with } x < 2^{\rho_m-1} \text{ then } \text{LWE}_s(1) \text{ else } \text{LWE}_s(0) */$
- 7 $\mathbf{ct}_{tmp} \leftarrow \mathbf{ct}_{tmp_m} + \mathbf{ct}_{tmp_e}; \quad /* \text{tmp equals zero only if } |\text{Dec}(\mathbf{ct}_f)| > 1 */$
 $/* \text{encrypt 0 if tmp == 0, 1 otherwise} */$
- 8 $\overline{\overline{\mathbf{CT}}}_{tmp} \leftarrow \text{CBS}(\mathbf{ct}_{tmp}, \text{PUB})$
- 9 $\mathbf{ct}_s \leftarrow \text{ExtendedCMux}(\mathbf{ct}_s, \mathbf{ct}_{s_{tmp}}, \overline{\overline{\mathbf{CT}}}_{tmp}),$
 $\mathbf{ct}_e \leftarrow \text{ExtendedCMux}(\mathbf{ct}_e, \mathbf{ct}_{e_{tmp}}, \overline{\overline{\mathbf{CT}}}_{tmp}),$
 $\mathbf{ct}_m \leftarrow \text{ExtendedCMux}(\mathbf{ct}_m, \mathbf{ct}_{m_{tmp}}, \overline{\overline{\mathbf{CT}}}_{tmp})$
- 10 **return** $\mathbf{ct}_f = [\mathbf{ct}_s, \mathbf{ct}_m, \mathbf{ct}_e]$

Future works

In future works, we aim to integrate the strengths of both methodologies presented in this paper: the minifloat and the HFP approaches. As detailed in Section 3.1, the WoP-PBS faces limitations when evaluating bivariate functions due to the exponential growth of the LUT. However, this constraint is less significant when WoP-PBS is used solely for computing univariate functions. Therefore, for intermediate precision (around 16 bits), we could adopt the representation proposed in Section 3.2, utilizing all algorithms introduced for HFP in combination with WoP-PBS where it proves efficient.

Additionally, recent improvements to the CBS in [WWL⁺24] should immediately enhance the running time of the algorithms introduced in this paper.

Although there is still a long way to go before achieving a fully practical deployment, this work narrows the gap between plaintext and homomorphic floating-point computations. This could open new avenues for implementing confidential machine learning algorithms, which often require a wide range of values.

References

- [AN16] Seiko Arita and Shota Nakasato. Fully homomorphic encryption for point numbers. In *International Conference on Information Security and Cryptology*, pages 253–270. Springer, 2016.
- [BBB⁺22] Loris Bergerat, Anas Boudi, Quentin Bourgerie, Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Parameter optimization & larger precision for (T)FHE. *IACR Cryptol. ePrint Arch.*, page 704, 2022.
- [BCL⁺23] Loris Bergerat, Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, Adeline Roux-Langlois, and Samuel Tap. Faster secret keys for (t) fhe. *Cryptology ePrint Archive*, 2023.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 309–325, 2012.
- [BIP⁺22] Charlotte Bonte, Ilia Iliashenko, Jeongeun Park, Hilder VL Pereira, and Nigel P Smart. Final: Faster fhe instantiated with ntru and lwe. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 188–215. Springer, 2022.
- [Bra12] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical gapsvp. *IACR Cryptology ePrint Archive*, 2012:78, 2012.
- [BST20] Florian Bourse, Olivier Sanders, and Jacques Traoré. Improved secure integer comparison via homomorphic encryption. In *Cryptographers’ Track at the RSA Conference*, pages 391–416. Springer, 2020.
- [CGGI20] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, 2020.

- [CHK⁺18] Jung Hee Cheon, Kyoohyung Han, Andrey Kim, Miran Kim, and Yongsoo Song. A full rns variant of approximate homomorphic encryption. In *International Conference on Selected Areas in Cryptography*, pages 347–368. Springer, 2018.
- [CJL⁺20] Ilaria Chillotti, Marc Joye, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Concrete: Concrete operates on ciphertexts rapidly by extending TfhE. In *WAHC 2020*, 2020.
- [CJP21] Ilaria Chillotti, Marc Joye, and Pascal Paillier. Programmable bootstrapping enables efficient homomorphic inference of deep neural networks. In *CSCML 2021*. Springer, 2021.
- [CKKS17] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. Homomorphic encryption for arithmetic of approximate numbers. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I*, pages 409–437, 2017.
- [CLOT21] Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Improved programmable bootstrapping with larger precision and efficient arithmetic circuits for tfhe. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 670–699. Springer, 2021.
- [CSVW16] Anamaria Costache, Nigel P Smart, Srinivas Vivek, and Adrian Waller. Fixed-point arithmetic in she schemes. In *International Conference on Selected Areas in Cryptography*, pages 401–422. Springer, 2016.
- [CZB⁺22] Pierre-Emmanuel Clet, Martin Zuber, Aymen Boudguiga, Renaud Sirdey, and Cédric Gouy-Pailler. Putting up the swiss army knife of homomorphic calculations by means of tfhe functional bootstrapping. Cryptology ePrint Archive, Report 2022/149, 2022. <https://ia.cr/2022/149>.
- [DM15] Léo Ducas and Daniele Micciancio. FHEW: bootstrapping homomorphic encryption in less than a second. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International*

- Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, pages 617–640, 2015.
- [FV12] Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. *IACR Cryptology ePrint Archive*, 2012:144, 2012.
- [GBA21] Antonio Guimarães, Edson Borin, and Diego F. Aranha. Revisiting the functional bootstrap in TFHE. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2021(2):229–253, 2021.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 169–178, 2009.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. *IACR Cryptology ePrint Archive*, 2013:340, 2013.
- [Hwa24] Vincent Hwang. Formal verification of emulated floating-point arithmetic in falcon. In *International Workshop on Security*, pages 125–141. Springer, 2024.
- [Ins08] Institute of Electrical and Electronics Engineers. Ieee standard for floating-point arithmetic. *IEEE Std 754-2008*, pages 1–70, 2008.
- [KO22] Jakub Klemsa and Melek Onen. Parallel operations over tfhe-encrypted multi-digit integers. *Cryptology ePrint Archive*, Report 2022/067, 2022.
- [Lai17] Kim Laine. Simple encrypted arithmetic library 2.3.1. *Microsoft Research* <https://www.microsoft.com/en-us/research/uploads/prod/2017/11/sealmanual-2-3-1.pdf>, 2017.
- [LMK⁺23] Yongwoo Lee, Daniele Micciancio, Andrey Kim, Rakyong Choi, Maxim Deryabin, Jieun Eom, and Donghoon Yoo. Efficient

- flew bootstrapping with small evaluation keys, and applications to threshold homomorphic encryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 227–256. Springer, 2023.
- [LMP21] Zeyu Liu, Daniele Micciancio, and Yuriy Polyakov. Large-precision homomorphic sign evaluation using flew/tfhe bootstrapping. Cryptology ePrint Archive, Report 2021/1337, 2021. <https://ia.cr/2021/1337>.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT 2010*. Springer, 2010.
- [LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3):565–599, 2015.
- [LS22] Seunghwan Lee and Dong-Joon Shin. Overflow-detectable floating-point fully homomorphic encryption. *Cryptology ePrint Archive*, 2022.
- [MBDD⁺18] Jean-Michel Muller, Nicolas Brisebarre, Florent De Dinechin, Claude-Pierre Jeannerod, Vincent Lefevre, Guillaume Melquiond, Nathalie Revol, Damien Stehlé, Serge Torres, et al. *Handbook of floating-point arithmetic*. Springer, 2018.
- [ML20] Subin Moon and Younho Lee. An efficient encrypted floating-point representation using heaan and tfhe. *Security and Communication Networks*, 2020, 2020.
- [PFH⁺20] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon. *Post-Quantum Cryptography Project of NIST*, 2020.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC 2005*. ACM, 2005.

- [SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT 2009*. Springer, 2009.
- [WK19] Shibo Wang and Pankaj Kanwar. Bfloat16: The secret to high performance on cloud tpus. *Google Cloud Blog*, 4, 2019.
- [WWL⁺24] Ruida Wang, Yundi Wen, Zhihao Li, Xianhui Lu, Benqiang Wei, Kun Liu, and Kunpeng Wang. Circuit bootstrapping: faster and smaller. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 342–372. Springer, 2024.
- [Zam22] 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>.