

Improved Circuit Synthesis with Multi-Value Bootstrapping for FHEW-like Schemes

Johannes Mono¹, Kamil Kluczniak² and Tim Güneysu^{1,3}

¹ Ruhr University Bochum, first.last@rub.de

² secunet Security Networks, kamil.kluczniak@gmail.com

³ DFKI GmbH

Abstract. In recent years, the research community has made great progress in improving techniques for privacy-preserving computation, such as fully homomorphic encryption (FHE). Despite the progress, there remain open challenges, mainly in performance and usability, to further advance the adoption of these technologies. This work provides multiple contributions that improve the current state-of-the-art in both areas. More specifically, we significantly simplify the multi-value bootstrapping by Carrov, Izabachène, and Mollimard [CIM19] for Boolean-based FHE schemes such as FHEW or TFHE, making the concept usable in practice. Based on our simplifications, we implement an easy-to-use interface for multi-value bootstrapping in the open-source library FHE-Deck [fhe23], derive new parameter sets for multi-bit encryptions with state-of-the-art security, and build a toolset that translates high-level code to multi-bit operations on encrypted data using circuit synthesis. We propose and integrate the first non-trivial FHE-specific optimizations for privacy-preserving circuit synthesis: look-up table (LUT) grouping and adder substitution. Using LUT grouping, we reduce the number of bootstrapping operations by almost 40% on average, while for adder substitution, we reduce the number of required bootstrappings by up to 85% for certain use cases. Overall, the execution time is up to $4.2\times$ faster with all optimizations enabled compared to previous state-of-the-art circuit synthesis.

Keywords: fully homomorphic encryption · FHEW · TFHE · circuit synthesis

1 Introduction

Encryption is a fundamental technology of today’s society, securing data and communications around the globe. One exciting application is privacy-preserving computation, where techniques such as fully homomorphic encryption (FHE) encrypt and protect sensitive data during computation. However, there remain open challenges in adopting these technologies, for example, reducing computational costs or improving usability.

Gentry’s seminal work [Gen09] introduces bootstrapping, in which the error associated with a homomorphic ciphertext is refreshed to allow for indefinite computation. Current state-of-the-art schemes are still based on this ingenious idea, and two strains have emerged. Word-based schemes such as BFV [Bra12, FV12], BGV [BGV14], and CKKS [CKKS17] operate on multiple large elements at a time and excel at highly parallelizable tasks. These schemes usually perform many operations followed by an expensive bootstrapping procedure and are often used for specific use cases requiring mainly additions or multiplications.

In contrast, Boolean-based schemes such as FHEW [DM15] and TFHE [CGGI16a] provide high flexibility encrypting single bits or small bit groups and are thus a good fit for a wide variety of use cases. A target function is commonly represented as a circuit composed of Boolean gates with encrypted input bits, and every gate evaluation requires a comparatively fast bootstrapping. However, bootstrapping is still the most expensive part

of such circuits, and thus, it is crucial to reduce the number of gates when translating a use case to a circuit.

Translating use cases is either done manually for critical tasks (this can be compared to hand-written assembly) or using tools to automate translation from high-level code to Boolean circuits (similar to code compilation). Although the former can result in the best performance for a given use case, the process is rather tedious, and there has been some effort by the research community to provide automatic translations from high-level code to circuits [CDS15, CMG⁺18, GSPH⁺21], improving usability and easing adoption for non-experts interested in FHE.

Currently, two different approaches exist for automatically converting high-level code to Boolean circuits. The first approach is based on instruction mapping, where the high-level code is translated to an intermediate representation (IR). Afterward, the individual instructions of the IR are mapped to Boolean primitives and composed accordingly. The second approach is based on existing hardware tooling using so-called synthesizers. First, high-level synthesis converts high-level code to a hardware description language (HDL). Then, a synthesizer converts the HDL to a Boolean circuit. Finally, no matter the approach, the resulting circuit is translated to FHE library code.

In this work, we introduce several contributions to enhance the current state-of-the-art in FHE transpilation, with a primary focus on the synthesis-based approach:

- We significantly simplify the multi-value bootstrapping idea proposed by Carpov, Izabachène, and Mollimard [CIM19], improving the current state-of-the-art for bootstrapping in FHEW-like schemes. Our modifications make it feasible to use this technique in practice and enhance the capabilities of homomorphic gates (Subsection 3.2).
- We adapt an optimization technique for single- and multi-bit circuits to the circuit synthesis process and handle additions specifically optimized for FHE circuits. This results in circuits requiring up to 85 % less bootstrappings compared to non-optimized circuits (Subsection 3.7).
- We perform novel post-synthesis optimizations on the netlist based on our improved multi-value bootstrapping to evaluate multiple gates at once. With this optimization, execution times are up to 2.4× faster compared to the previous state-of-the-art (Subsection 3.6).
- We propose new parameter sets for FHEW-like encryptions with state-of-the-art security for multi-value bootstrapping. We also implement an easy-to-use interface for multi-value bootstrapping in the open-source library FHE-Deck [fhe23] (Subsection 3.4).

2 Preliminaries

In the following, we introduce the background to our work. We start with notations used throughout the paper, provide a formal definition of Boolean and homomorphic circuits, shortly describe FHE and its hardness assumption and conclude with a short section on transpilation in the context of FHE.

2.1 Notation

We denote as \mathbb{Z}_Q the group of integers modulo Q and, for a power-of-two N , as \mathcal{R}_Q the ring of polynomials $\mathbb{Z}_Q[X]/(X^N + 1)$. We call $z \in \mathcal{R}_Q$ with $z^n = 1 \in \mathcal{R}_Q$ the n -th root of unity. Note that \mathcal{R}_Q has $2N$ roots of unity of the form X^a for $a \in \mathbb{Z}_{2N}$. Furthermore, the roots of unity in \mathcal{R}_Q form an algebraic group of order $2N$ with respect to multiplication. For

clarity, we denote ring elements using the `\frac` environment e.g. $\mathbf{a} \in \mathcal{R}_Q$. We denote a n -dimensional column vector as $[f(\cdot, i)]_{i=1}^n$, where $f(\cdot, i)$ defines the i -th coordinate. For brevity, we will also denote as $[n]$ the vector $[i]_{i=1}^n$, and as $[n, m]_{i=n}^m$ the vector $[n, \dots, m]^\top$. We address the i -th entry of a vector \vec{v} by $\vec{v}[i]$. We denote the inner product between vector \vec{a} and \vec{b} as $\langle \vec{a}, \vec{b} \rangle$.

By $x \leftarrow_R \mathcal{S}$, we denote sampling a random variable from the set \mathcal{S} . By default, we sample from the uniform distribution and explicitly state when referring to other distributions. For a random variable $a \in \mathbb{Z}$, we denote as $\text{Var}(a)$ the variance of a , its expectation as $\mathbb{E}(x)$, and its standard deviation as $\text{SD}(a)$. For $\mathbf{a} \in \mathcal{R}_Q$, we define $\text{Var}(\mathbf{a})$ and $\mathbb{E}(\mathbf{a})$ to be the variance and expectation of the coefficients of the polynomial \mathbf{a} , respectively. We denote any polynomial as $\text{poly}(\cdot)$. We denote as $\text{negl}(\lambda)$ a negligible function in $\lambda \in \mathbb{N}$; that is, for any positive polynomial $\text{poly}(\cdot)$ there exists $c \in \mathbb{N}$ such that for all $\lambda \geq c$ we have $\text{negl}(\lambda) \leq \frac{1}{\text{poly}(\lambda)}$.

2.2 Boolean and Homomorphic Circuits

We define a Boolean gate as an arbitrary Boolean function $\mathcal{G}_n : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2$ for a positive integer n . Any Boolean gate can be represented as LUT which stores 2^n outputs $O \in \mathbb{F}_2$, one for each input $I \in \mathbb{Z}_2^n$; we often use the terms Boolean gate and LUT interchangeably. An example for a Boolean gate \mathcal{G}_1 is an inverter gate

$$\text{INV} : \mathbb{Z}_2 \rightarrow \mathbb{Z}_2, \quad I \mapsto O = I + 1.$$

We can generalize a Boolean gate to multiple outputs as $\mathcal{G}_{n,m} : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2^m$ which can be represented with m LUTs, one for each output bit. Here, an example is a full adder with the inputs $x, y \in \mathbb{Z}_2$ and a carry-in $c_i \in \mathbb{Z}_2$ defined as

$$\text{FA} : \mathbb{Z}_2^3 \rightarrow \mathbb{Z}_2^2, \quad (x, y, c_i) \mapsto (s, c_o) = (x + y + c_i, x \cdot y + c_i \cdot (x + y))$$

for the sum s and the carry-out c_o . The number of input wires is also called fan-in of a gate, and the number of output wires fan-out.

A Boolean circuit $\mathcal{C}_2 : \mathbb{Z}_2^s \rightarrow \mathbb{Z}_2^t$ is a directed acyclic graph where the vertices are Boolean gates $\mathcal{G}_{n,m}$ and the edges connect gate outputs to the inputs of other gates with s unconnected global inputs and t unconnected global outputs. A well-known fact is that, for every directed acyclic graph, there exists a topological ordering of the graph. A topological ordering is an ordering such that for every edge, the start vertex of this edge appears before the end vertex in the sorted list of vertices. For a circuit \mathcal{C}_2 , sorting the graph topologically and evaluating the sorted gates for a global input $I \in \mathbb{Z}_2^s$ ensures that, for any given gate $\mathcal{G}_{n,m}$, all n inputs are known.

We can generate a Boolean circuit \mathcal{C}_2 from a circuit design in a HDL through a process called synthesis. The output, a textual representation of \mathcal{C}_2 , is also referred to as netlist. Generally, a netlist can include additional elements such as registers for storage; however, in this work, a netlist only contains Boolean gates. During synthesis, the circuit is optimized, mostly heuristically, according to specific parameters, such as area usage or power consumption. Extending the idea of synthesis to high-level code is called high-level synthesis. Examples are the high-level synthesis tool XLS¹, translating C++ code to the HDL Verilog, and the low-level synthesis tool Yosys², generating a netlist from Verilog.

Usually, synthesis tools output a netlist either based on a gate library, a list of available gates for a given hardware platform, or output a LUT-based netlist with many LUTs for different functions \mathcal{G}_n and some upper bound on fan-in³. For hardware with special high-

¹<https://github.com/google/xls>

²<https://github.com/yosyshq/yosys>

³We note for completeness that synthesis tools also place gates $\mathcal{G}_{n,m}$ with $m > 1$, such as registers. Since we configure these tools to only output LUTs, which, by definition, only have a single output, this work assumes a netlist only containing a Boolean circuit \mathcal{C}_2 composed of gates \mathcal{G}_n .

performance gates, state-of-the-art synthesis tools are able to replace certain subcircuit patterns with optimized primitives; for example, multiplications are often realized in digital signal processing units.

We define a homomorphic circuit $\mathcal{C} : \mathbb{Z}_p^s \rightarrow \mathbb{Z}_p^t$ consisting of homomorphic gates of the form $f(b + \sum x_i \cdot a_i)$ for a function $f : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$, known scalars $a_i, b \in \mathbb{Z}_p$, and encrypted plaintexts x_i . We compute the affine part $b + \sum x_i \cdot a_i$ using the homomorphic capabilities of the FHE scheme and the function f during bootstrapping, also known as functional bootstrapping (see Subsection 3.2). Note that we compute the affine part over the integers, but require the input to f in \mathbb{Z}_p ; this has to be taken into account during homomorphic circuit design.

2.3 Fully Homomorphic Encryption

An FHE scheme consists of four algorithms (Setup, Enc, Eval, Dec), each with the following syntax [RAD78, Gen09].

- **Setup**(λ): This probabilistic polynomial time (PPT) algorithm takes as input a security parameter λ and outputs an evaluation key ek and a secret key sk .
- **Enc**(sk, m): This PPT algorithm takes as input a secret key sk as well as a message m and returns a ciphertext ct .
- **Eval**($\text{ek}, [\text{ct}_i]_{i=1}^n, \mathcal{C}$): Given an evaluation key ek , ciphertexts $[\text{ct}_i]_{i=1}^n$, and a circuit \mathcal{C} , this (non-)deterministic algorithm outputs a ciphertext ct .
- **Dec**(sk, ct): Given a secret key sk and a ciphertext ct , this deterministic algorithm outputs a message m .

Informally, we say that an FHE scheme is correct, if the outcome of the evaluation of a circuit \mathcal{C} on ciphertexts encrypting messages $\text{m}_1, \dots, \text{m}_n$ decrypts to $\mathcal{C}(\text{m}_1, \dots, \text{m}_n)$. Formally, we say that FHE is correct if for all security parameters $\lambda \in \mathbb{N}$, the circuits $\mathcal{C} : \mathcal{M}^n \rightarrow \mathcal{M}$ over the message space \mathcal{M} of depth $\text{poly}(\lambda)$, and all messages $[\text{m}_i \in \mathcal{M}]_{i=1}^n$ we have

$$\Pr [\text{Dec}(\text{sk}, \text{ct}_{\text{out}}) = \mathcal{C}([\text{m}_i]_{i=1}^n)] = 1 - \text{negl}(\lambda),$$

where $(\text{ek}, \text{sk}) \leftarrow \text{Setup}(\lambda)$, $\text{Dec}(\text{sk}, \text{ct}_i) = \text{m}_i$ for all $i \in [n]$ and $\text{ct}_{\text{out}} \leftarrow \text{Eval}(\text{ek}, [\text{ct}_i]_{i=1}^n, \mathcal{C})$. For efficiency, we require that Setup, Enc and Dec run in polynomial time in the security parameter, that is $\text{poly}(\lambda)$, and Eval runs in $\text{poly}(\lambda, |\mathcal{C}|)$. Finally, we say that an FHE scheme is compact if the size of the output of Eval is independent of the size of the circuit \mathcal{C} . More specifically, we require that $|\text{Eval}(\text{ek}, [\text{ct}_i]_{i=1}^n, \mathcal{C})|$ is $\text{poly}(\lambda, |\mathcal{M}|)$.

Indistinguishability Under Chosen Plaintext Attack. Let $\lambda \in \mathbb{N}$ be a security parameter and $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ be a PPT adversary. We say that an FHE scheme is indistinguishable under chosen plaintext attack (IND-CPA)-secure if the probability

$$\Pr \left[\begin{array}{l} \text{sk} \leftarrow \text{Setup}(\lambda), \\ (\text{st}, \text{m}_0, \text{m}_1) \leftarrow \mathcal{A}_0^{\mathcal{O}(\text{sk}, \cdot)}(\lambda), \\ b \leftarrow_R \{0, 1\}, \\ \text{ct}_b \leftarrow \text{Enc}(\lambda, \text{sk}, \text{m}_b) \end{array} \right] \mathcal{A}_1(\text{ct}_b, \text{st}) = b:$$

is at most $\text{negl}(\lambda)$ for all PPT adversaries \mathcal{A} ; the oracle \mathcal{O} on input of a message m outputs $\text{ct} \leftarrow \text{Enc}(\text{sk}, \text{m})$.

2.4 Generalized Learning with Errors

Definition 1 (GLWE). Let \mathcal{D}_{sk} be a (not necessarily uniform) distribution over \mathcal{R}_Q , and $\mathcal{D}_{\mathcal{R},\sigma}$ be a noise distribution over \mathcal{R}_Q with standard deviation $\sigma > 0$, $n \in \mathbb{N}$ and $N \in \mathbb{N}$ be a power-of-two, that are chosen according to a security parameter λ . For $\vec{\mathbf{a}} \leftarrow_R \mathcal{R}_Q^n$, $\mathbf{e} \leftarrow_R \mathcal{D}_{\mathcal{R},\sigma}$ and $\vec{\mathbf{s}} \in \mathcal{D}_{\text{sk}}^n$, we define a Generalized Learning with Errors (GLWE) sample of a message $\mathbf{m} \in \mathcal{R}_Q$ with respect to $\vec{\mathbf{s}}$, as

$$\text{GLWE}_{\sigma,n,N,Q}(\vec{\mathbf{s}}, \mathbf{m}) = \begin{bmatrix} -\vec{\mathbf{a}}^\top \cdot \vec{\mathbf{s}} + \mathbf{e} \\ \vec{\mathbf{a}}^\top \end{bmatrix} + \begin{bmatrix} \mathbf{m} \\ \vec{\mathbf{0}} \end{bmatrix} \in \mathcal{R}_Q^{(n+1)}.$$

We say that the $\text{GLWE}_{\sigma,n,N,Q}$ -assumption holds if for any PPT adversary \mathcal{A} we have

$$\left| \Pr [\mathcal{A}(\text{GLWE}_{\sigma,n,N,Q}(\vec{\mathbf{s}}, 0))] - \Pr [\mathcal{A}(\mathcal{U}_Q^{n+1})] \right| \leq \text{negl}(\lambda)$$

where \mathcal{U}_Q^{n+1} is the uniform distribution over \mathcal{R}_Q^{n+1} .

We denote a Learning with Errors (LWE) sample as $\text{LWE}_{\sigma,n,Q}(\vec{\mathbf{s}}, m) = \text{GLWE}_{\sigma,n,1,Q}(\vec{\mathbf{s}}, m)$, which is a special case of a GLWE sample where the ring is $\mathbb{Z}_Q[X]/(X + 1)$. Note that, in this case, $\vec{\mathbf{s}} \in \mathbb{Z}_p$ is a integer vector and $m \in \mathbb{Z}_p$. Similarly, we denote a Learning with Errors over Rings (RLWE) sample as $\text{RLWE}_{\sigma}(\mathbf{s}, \mathbf{m}) = \text{GLWE}_{\sigma,1,N,Q}(\mathbf{s}, \mathbf{m})$ which is the special case of an GLWE sample with $n = 1$. For simplicity, we omit to state the modulus and ring dimension for RLWE samples because we always use $\mathcal{R}_Q = \mathbb{Z}_Q[X]/(X^N + 1)$ where N is a power-of-two. For LWE samples, we will be switching between different moduli and different dimensions; hence we will indicate the current modulus in the notation. We sometimes leave the inputs unspecified and substitute them with a dot (\cdot) when it is not necessary to refer to them within the scope of a function. We define the error of $\vec{\mathbf{c}} \in \mathcal{R}_Q^{(n+1)}$ as $\text{Error}(\vec{\mathbf{c}}, \mathbf{m}) = \langle \vec{\mathbf{c}}, \vec{\mathbf{s}} \rangle - \mathbf{m}$. Finally, we define the symbol $\Delta_{Q,p} = \lfloor \frac{Q}{p} \rfloor$ where $p \in \mathbb{N}$ is the plaintext modulus.

2.5 FHE Transpilation

Transpilation, also known as source-to-source compilation, is a process in which the source code written in one programming language is converted into the source code of another language. It converts code at similar levels of abstraction without changing the code’s logic or functionality. In the context of FHE, transpilation mostly refers to converting high-level code implementing functionality in the unencrypted domain to FHE library code in the homomorphic realm. For example, the FHE transpiler [GSPH⁺21] converts a subset of C++ code to C++ or Rust, depending on the chosen output library.

Although FHE transpilation operates at a similar abstraction level with respect to the input and output programming language, the process itself closely resembles a compilation process as FHE libraries commonly only implement low-level operations on the encrypted data. Thus, FHE transpilation first converts high-level code to an IR where optimizations are performed. Then, the IR is further processed with instruction mapping or synthesis. For instruction mapping, each IR instruction is mapped to the low-level operations exposed by the FHE library, while for synthesis, hardware synthesis tools are used to convert the IR to a low-level circuit matching the low-level operations provided by the chosen library.

3 Contributions

We now describe our theoretical contributions across the circuit synthesis process and start with the synthesis toolchain top to bottom. Then, we highlight our contributions in the reverse starting with the modified FHE bootstrapping and ending with the optimizations for the synthesis process.




Figure 1: Our synthesis toolchain with different stages and their transitions.

3.1 Synthesis Toolchain

Our synthesis toolchain closely follows the structure and even reuses parts of the FHE transpiler [GSPH⁺21], which currently provides the best publicly available solution to circuit synthesis. More specifically, high-level C++ code is converted to Verilog via the XLS framework⁴. Then, as in the transpiler, we synthesize Verilog to a Boolean circuit \mathcal{C}_2 via Yosys⁵. We configure synthesis in Yosys with pre-defined Yosys scripts specifying the exact steps Yosys should execute during synthesis. Instead of using the rather simple script of the transpiler, we create our own scripts based on the default Yosys script for LUT-based hardware. In contrast to the transpiler, there are two important differences:

- Using our new synthesis scripts, we can substitute certain subcircuits with custom gates. We use this to replace additions and subtractions with an optimized homomorphic gate structure using FHE-specific primitives (Subsection 3.7).
- We perform post-synthesis optimizations on the circuit \mathcal{C}_2 with HAL [Emb19], a netlist analysis tool, using its Python interface (Subsection 3.6, Subsection 3.8). Our Python scripts directly translates the optimized circuit to FHE-Deck code [fhe23] mapping the circuit as described in Subsection 3.5.

In Figure 1, we depict the different stages of our toolchain including their transitions. All of our code, including tests and examples, is publicly available on GitHub⁶.

3.2 Efficient Multi-Value Functional Bootstrapping

In the following, we describe our multi-value bootstrapping algorithm, our instantiation of the FHE scheme and how we enhance homomorphic gate evaluation from the function f to a function $f_m : \mathbb{Z}_p \rightarrow \mathbb{Z}_p^m$ at the cost of computing just one f . Since the multi-value bootstrapping algorithm is based on FHEW-style bootstrapping algorithms [DM15, CGGI16a], we first give an informal overview of these algorithms in a black-box fashion limiting ourselves to their input-output relation, which is enough to understand the bootstrapping algorithm.

Among the most efficient types of FHE schemes are schemes based on the FHEW-style bootstrapping algorithms [DM15]. There are multiple variants [CGGI16a, MS18, CIM19, BDF18, GBA21, Klu22b] and improvements [YXS⁺21, CLOT21, LMP22, KS22] of these algorithms. The goal is to homomorphically compute the decryption function on an LWE ciphertext and compute an arbitrary function on the encrypted plaintext message along the way.

Suppose we have an LWE ciphertext \vec{c} with $\langle \vec{c}, \vec{s} \rangle = M + e$, where \vec{s} is the secret key, M the message, and e the error. The operation $\langle \vec{c}, \vec{s} \rangle$ can be realized within a cyclic algebraic group, more specifically, the group of rotations. The idea is to realize the rounding

⁴<https://github.com/google/xls>

⁵<https://github.com/yosyshq/yosys>

⁶<https://github.com/Chair-for-Security-Engineering/fhewsyn>

function $\lfloor \cdot \rfloor$ by setting the elements of the vector such that messages are encoded in intervals of appropriate size to handle the noise term e . Bootstrapping algorithms for FHEW-like schemes use the design pattern established by Alperin-Sheriff and Peikert [AP14] over polynomial rings. In particular, the observation first made by Ducas and Micciancio [DM15] is that in the ring $\mathbb{Z}[X]/(X^N + 1)$, the product of any ring element with a root of unity (negacyclicly) rotates the coefficients of that ring element. In other words, given a polynomial $\mathbf{w} = \sum_{i=0}^{N-1} w_i \cdot X^i$, we have

$$\mathbf{w} \cdot X^y = \sum_{i=0}^{N-y-1} w_i \cdot X^{i+y} - \sum_{i=N-y}^{N-1} w_i \cdot X^{i+y-N}.$$

As part of the blind rotation procedure, we compute $\mathbf{w} \cdot X^{(\vec{c}, \vec{s})} = \mathbf{w} \cdot X^{M+e}$ homomorphically. Since computation takes place over RLWE ciphertexts, we obtain at the end of the blind rotation procedure an RLWE ciphertext of $\mathbf{w} \cdot X^{M+e}$. Finally, Ducas and Micciancio [DM15] observe that given such RLWE ciphertext, one can extract an LWE ciphertext that encrypts the constant coefficient of the message; more specifically, the element $\mathbf{w} \cdot X^{M+e}[1]$. The step is done via the `SampleExtract` procedure (see Table 1). Then, the final step is to choose the polynomial \mathbf{w} such that $\mathbf{w} \cdot X^{M+e}[1]$ encodes the desired value, according to a given function $f : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$, and switch the extracted LWE ciphertext to an LWE ciphertext that is suitable for another bootstrapping step.

It is easy to see that with a bootstrapping algorithm which computes a function $f : \mathbb{Z}_p \rightarrow \mathbb{Z}_p$, we can modify it to compute a function $f_m : \mathbb{Z}_p \rightarrow \mathbb{Z}_p^m$ by grouping m distinct f , one for every output of f_m ; this would require m bootstrapping invocations. The goal of multi-value bootstrapping is to compute all m functions at the cost of only one bootstrapping. The general idea to amortize computation of different functions on the same input ciphertext is based on a previous work by Carпов, Izabachène, and Mollinard [CIM19], but in this paper, we significantly simplify execution of the idea. We give a version of the bootstrapping algorithm Algorithm 1 and provide a bound on the bootstrapping noise in Theorem 1.

The algorithm takes as input the blind rotation key `brKey` and the key switching key `ksKey`, an LWE ciphertext $\vec{c} \in \mathbb{Z}_Q^{N+1}$, and polynomials \mathbf{w} and $\vec{\mathbf{v}}$. First, the algorithm switches the LWE key; this results in an LWE ciphertext \vec{c}_{ksKey} with a smaller dimension $n \in \mathbb{N}$ and a secret key from a smaller distribution, for instance binary, ternary, or Gaussian. Afterward, the algorithm switches the modulus from Q to $2N$ (recall that the roots of unity in the ring \mathcal{R}_Q form an algebraic group of order $2N$). Then, we run `BlindRotate` which homomorphically computes $\mathbf{m}_{\text{acc}} \leftarrow \mathbf{w} \cdot X^{(\vec{c}, \vec{s})}$. Finally, we execute a for-loop that multiplies the ciphertext with a polynomial from the vector $\vec{\mathbf{v}}$. Consequently, we obtain and return a vector of ciphertexts $[\vec{c}_{\text{out}, i}]_{i=1}^k$. The i -th ciphertext in the returned vector encrypts the message $\mathbf{m}_{\text{acc}} \cdot \vec{\mathbf{v}}[i]$. The problem when using this construction is that the multiplications by the elements of the vector $\vec{\mathbf{v}}$ may blow up the error, ultimately destroying the ciphertext. This can happen if the norm bounds of elements in $\vec{\mathbf{v}}$ are too large.

In our work, we provide a simple solution to build the polynomials \mathbf{w} and $\vec{\mathbf{v}}$ with time linear in N . We observe that when we are only interested in extracting bits the polynomials in $\vec{\mathbf{v}}$ are already sparse and of small infinity norm. Later, we compute a simple binary composition before running the new bootstrapping algorithm and computing the next LUT function. Concretely, we set $\mathbf{w} = \Delta_{Q,p}$. Then, for $i \in [k]$, the polynomial $\vec{\mathbf{v}}[i]$ encoding the function $f^{(i)} : \mathbb{Z}_p \rightarrow \mathbb{Z}_2$ is defined as

$$\vec{\mathbf{v}}[i] = f^{(i)}(0) - \sum_{j=1}^{N-1} f^{(i)}\left(\lfloor \frac{p}{2N} \cdot j \rfloor\right) \cdot X^{N-j}. \tag{1}$$

In other words, for each $i \in [k]$, we set the constant coefficient $\mathbf{v}[i][1]$ to $f^{(i)}(0)$, and $\mathbf{v}[i][N - j]$ to $-f^{(i)}(\lfloor j/2N \rfloor)$ for $j \in [N - 1]$. Note that the construction of $\mathbf{v}[i]$ involves

Algorithm 1: Bootstrap(brKey, ksKey, \vec{c} , \mathfrak{w} , $\vec{\mathfrak{v}}$)**Input:**

The blind rotation key brKey;
 The key switching key ksKey;
 An LWE ciphertext $\vec{c} \in \mathbb{Z}_Q^{N+1}$;
 Polynomial $\mathfrak{w} \in \mathcal{R}_Q$; and
 A vector of polynomials $\vec{\mathfrak{v}} \in \mathcal{R}_p^k$.

Output: We consider two different versions of the bootstrapping algorithm: the classic variant and the multi-value variant. For the classic variant, the vector $\vec{\mathfrak{v}}$ is empty ($k = 0$), and the algorithm returns a ciphertext $\text{LWE}_{\cdot, N, Q}(\vec{s}, \mathfrak{w} \cdot X^{(\vec{c}, \vec{s})})$.

For the multi-value variant, the algorithm returns a vector of LWE ciphertexts $[\vec{c}_{\text{out}, i}]_{i=1}^k$, where $\vec{c}_{\text{out}, i} = \text{LWE}_{\cdot, N, Q}(\vec{s}, \mathfrak{w} \cdot X^{(\vec{c}, \vec{s})} \vec{\mathfrak{v}}[i])$.

begin

```

  Run  $\vec{c}_{\text{ksKey}} \leftarrow \text{KeySwitch}(\vec{c}, \text{ksKey}) \in \mathbb{Z}_Q^{n+1}$  ;
  Run  $\vec{c}_{\text{in}} \leftarrow \text{ModSwitch}(\vec{c}_{\text{ksKey}}, 2N) \in \mathbb{Z}_{2N}^{n+1}$  ;
   $\vec{c}_{\text{acc}} \leftarrow \text{BlindRotate}(\text{brKey}, \mathfrak{w}, \vec{c}_{\text{in}})$  ;
  for  $i = 1 \dots k$  do
    Compute  $\vec{c}_{\text{acc}, i} \leftarrow \vec{c}_{\text{acc}} \cdot \vec{\mathfrak{v}}[i]$  ;
    Compute  $\vec{c}_{\text{out}, i} \leftarrow \text{SampleExtract}(\vec{c}_{\text{acc}, i})$  ;
  Return  $[\vec{c}_{\text{out}, i}]_{i=1}^k$  ;

```

only N invocations of $f^{(i)}$. Therefore, the algorithm runs in time linear in N given that $f^{(i)}$ is independent of N .

Recall that we use functions that map an input from \mathbb{Z}_p to \mathbb{Z}_2 and need to keep the norm of the polynomial small. There is, however, no other reason for the output domain to be binary, and the algorithm can be generalized to a larger output domain, given that we compensate for a larger noise growth with the parameters. Note that the bootstrapping output is a vector of LWE ciphertexts encoding bits, but within the \mathbb{Z}_p plaintext space.

We now analyze the noise growth and correctness for the multi-value bootstrapping in the binary case. Hence, the infinity norm of the polynomials in $\vec{\mathfrak{v}}$ can be bounded by N .

Theorem 1 (Bootstrapping Correctness). *Let $\vec{c}_{\text{acc}} \in \mathcal{R}_Q^2$ be an RLWE ciphertext returned by BlindRotate in an execution of Algorithm 1. Let $\mathfrak{e}_{\text{acc}} = \text{Error}(\vec{c}_{\text{acc}}, \mathfrak{m}_{\text{acc}})$ where $\mathfrak{m}_{\text{acc}} = \mathfrak{w} \cdot X^M$, and $M = \langle \vec{c}_{\text{in}}, \vec{s} \rangle \bmod 2N$. Then, for all $i \in [k]$, we have*

$$\text{SD}(\text{Error}(\vec{c}_{\text{acc}, i}, \mathfrak{m}_{\text{acc}} \cdot \vec{\mathfrak{v}}[i])) \leq \sqrt{2N \cdot \text{Var}(\mathfrak{e}_{\text{acc}})}$$

Furthermore, if $\mathfrak{w} = \Delta_{Q,p}$ and $\vec{\mathfrak{v}}$ are chosen to encode functions $f^{(i)}$ as given by Equation 1, and if $\lfloor \frac{2N}{p} \cdot M \rfloor = m < p/2$, then $\mathfrak{m}_{\text{acc}} \cdot \vec{\mathfrak{v}}[i] = \Delta_{Q,p} \cdot f^{(i)}(m)$.

Proof. Recall that we assume that at most N of all coefficients in the polynomials in the $\vec{\mathfrak{v}}$ vector are non-zero and that $\mathfrak{e}_{\text{acc}} \in \mathcal{R}_Q$. When multiplying the RLWE ciphertext \vec{c}_{acc} by $\vec{\mathfrak{v}}[i]$, we multiply the resulting error polynomial which is then equal to $\mathfrak{e}_{\text{acc}, i} = \mathfrak{e}_{\text{acc}} \cdot \vec{\mathfrak{v}}[i]$. The d -th coefficient of $\mathfrak{e}_{\text{acc}, i}$ can be written as

$$\mathfrak{e}_{\text{acc}, i}[d] = \sum_{j=1}^d \mathfrak{e}_{\text{acc}}[j] \cdot \vec{\mathfrak{v}}[i][d - i + 1] + \sum_{j=d+1}^N \mathfrak{e}_{\text{acc}}[j] \cdot \vec{\mathfrak{v}}[i][N + d - i + 1].$$

Crucially, observe that the sum takes each coefficient from the polynomials only once and that at most N of the coefficients of $\vec{\mathfrak{v}}[i]$ are non-zero. All non-zero coefficients of $\vec{\mathfrak{v}}[i]$ are

either 1 or -1 . Hence, we have that

$$SD(\epsilon_{\text{acc},i}) \leq \sqrt{\sum_{i=1}^N \text{Var}(\epsilon_{\text{acc}})} \leq \sqrt{N \cdot \text{Var}(\epsilon_{\text{acc}})}.$$

What is left to show is the correctness of the output messages. For $i \in [k]$, the relevant output is given by the constant coefficient of

$$\mathbf{m}_{\text{acc}} \cdot \vec{v}[i] = \Delta_{Q,p} \cdot X^M \cdot \left(f^{(i)}(0) - \sum_{j=1}^{N-1} f^{(i)}\left(\lfloor \frac{p}{2N} \cdot j \rfloor\right) \cdot X^{N-j} \right). \tag{2}$$

For $M = 0$ we have $m = 0$ and the constant coefficient of the polynomial from Equation 2 is $\Delta_{Q,p} \cdot f^{(i)}(0)$. For $0 < M < N$ the constant coefficient is given by the index j that zeroes the exponent above X (that is, j such that $N - j + M \bmod N = 0$). Given $m < p/2$, we have that $M < N$, and the index j that points to the constant coefficient can only be equal to M . The value of the constant coefficient is set to $\Delta_{Q,p} \cdot f^{(i)}\left(\lfloor \frac{p}{2N} \cdot j \rfloor\right) = \Delta_{Q,p} \cdot f^{(i)}\left(\lfloor \frac{p}{2N} \cdot M \rfloor\right) = f^{(i)}(m)$. \square

3.3 The FHE Scheme

We now combine the algorithms from Table 1, containing subprocedures from previous work [DM15], and our new multi-value bootstrapping from Algorithm 1 to an FHE system fitting the formal scheme definition in Subsection 2.3 and our model of computation. We provide high-level interfaces in FHE-Deck [fhe23] for all algorithms, making them easily accesible for researchers and developers.

Setup(λ). The setup algorithm consists of three main parts.

1. We choose the modulus Q , a power-of-two dimension N of the ring \mathcal{R}_Q , and an LWE dimension $n \in \mathbb{N}$ according to the security parameter λ . Then, we choose $\mathfrak{s} \in \mathcal{R}_Q$ for the RLWE key and set \vec{s}_{ext} to be the coefficient vector of \mathfrak{s} . We choose $\vec{s} \in \{0, 1\}^n$ for the LWE key.
2. We run $\text{ksKey} \leftarrow \text{KSSetup}(\vec{s}, \vec{s}_{\text{ext}}, \ell_{\text{ksKey}}, \sigma_{\text{ksKey}})$.
3. We run $\text{brKey} \leftarrow \text{BRSetup}(\vec{s}, \mathfrak{s}, \ell_{\text{brKey}}, \sigma_{\text{brKey}})$.

Finally, we set the evaluation key $\text{ek} = (\text{brKey}, \text{ksKey})$ and the secret key $\text{sk} = (\mathfrak{s}, \vec{s}_{\text{ext}}, \vec{s})$.

Enc(sk, m). To encrypt a message $m' \in \mathbb{Z}_p$, we compute

$$\vec{c} \leftarrow \text{LWE}_{\sigma, N, Q}(\vec{s}_{\text{ext}}, m) \in \mathbb{Z}_Q^{N+1},$$

where $m = \frac{Q}{p} \cdot m' \in \mathbb{Z}_Q$.

Eval($\text{ek}, [\text{ct}_i]_{i=1}^n, \mathcal{C}$). For the homomorphic circuit \mathcal{C} , we compute each individual gate

$$f_m \left(\sum_{i=1}^k x_i \cdot 2^{i-1} \in \mathbb{Z}_p \right) \in \mathbb{Z}_p^m,$$

where $p = 2^k$ and k is the maximum fan-in of the gates in the circuit \mathcal{C} . Thus, we compute m values in \mathbb{Z}_p at an amortized cost of running just one bootstrapping from Algorithm 1. While we can output arbitrary values in \mathbb{Z}_p if accounted for in the parameters, we limit ourselves to output m bits in $\mathbb{Z}_2 \subseteq \mathbb{Z}_p$.

Table 1: List of subprocedures commonly used in FHEW-like schemes [DM15].

Key Switching	
KSSetup	<p>Input: Takes as input two LWE secret keys $\vec{s} \in \{0, 1\}^n$, $\vec{s}_{\text{ext}} \in \mathbb{Z}_Q^N$, a performance parameter $\ell_{\text{ksKey}} \in \mathbb{N}$, and a standard deviation $\sigma_{\text{ksKey}} \in \mathbb{R}$.</p> <p>Output: Generates a key switching key ksKey which consists of $N \cdot \ell_{\text{ksKey}}$ $\text{LWE}_{\sigma_{\text{ksKey}}, n, Q}(\vec{s}, \cdot)$ ciphertexts.</p>
KeySwitch	<p>Input: Takes as input a key switching key ksKey and an $\text{LWE}_{\cdot, N, Q}(\vec{s}_{\text{ext}}, \mathbf{m})$ sample of a message $\mathbf{m} \in \mathbb{Z}_Q$.</p> <p>Output: Returns an $\text{LWE}_{\cdot, n, Q}(\vec{s}, \mathbf{m})$ sample under the key \vec{s} encoding the same message \mathbf{m}.</p> <p>Description: The key switching process consists of $N \cdot \ell_{\text{ksKey}}$ scalar multiplications in \mathbb{Z}_Q. The parameter ℓ_{ksKey} largely determines the time and space efficiency; that is, the smaller ℓ_{ksKey}, the faster the computation and the smaller the space complexity of the key material, but the bigger the noise induced by the key switching operation.</p>
Blind Rotation	
BRSetup	<p>Input: Takes as input the LWE key $\vec{s} \in \{0, 1\}^n$, a RLWE key $\mathfrak{s} \in \mathcal{R}_Q$, a performance parameter $\ell_{\text{brKey}} \in \mathbb{N}$, and a standard deviation σ_{brKey}.</p> <p>Output: Generates a blind rotation key brKey that consists of $2n\ell_{\text{brKey}}$ $\text{RLWE}_{\sigma_{\text{brKey}}}(\mathfrak{s}, \cdot)$ ciphertexts.</p>
BlindRotate	<p>Input: Takes as input a blind rotation key brKey, an LWE sample ct under modulus $2N$, and an accumulator $\text{acc} = \text{RLWE}(\mathfrak{s}, \vec{m}_{\text{acc}})$.</p> <p>Output: BlindRotate returns a sample $\text{acc}_{\text{out}} = \text{RLWE}(\mathfrak{s}, \vec{m}_{\text{out}})$ with $\vec{m}_{\text{out}} = \vec{m}_{\text{acc}} \cdot X^{(\text{ct}, \vec{s})}$. The blind rotation process [DM15, CGGI16a] consists of $2n \cdot (\ell_{\text{ksKey}} + 1)$ polynomial multiplications of elements in \mathcal{R}_Q. In particular, at the heart of a blind rotation algorithm, there is a ring version of the GSW cryptosystem [GSW13]. In this paper, we consider the concrete blind rotation from [CGGI16a], hence the number of polynomial multiplications. Similarly to key switching, the smaller the parameter ℓ_{brKey}, the faster the blind rotation algorithms and the smaller the blind rotation key (at the cost of larger noise).</p>
Other	
ModSwitch	<p>Input: Takes as input an LWE sample $\text{LWE}_{\cdot, n, Q}(\vec{s}, \Delta_{Q,p} \cdot \mathbf{m})$ and a modulus $q < Q$ with $\mathbf{m} \in \mathbb{Z}_p$.</p> <p>Output: Returns an LWE sample $\text{LWE}_{\cdot, n, q}(\vec{s}, \Delta_{q,p} \cdot \mathbf{m})$ under modulus q.</p>
SampleExtract	<p>Input: Takes as input a RLWE encryption $\text{RLWE}_{\sigma_{\text{brKey}}}(\mathfrak{s}, \mathbf{m})$ of a message $\mathbf{m} \in \mathcal{R}_Q$.</p> <p>Output: Returns an LWE sample $\text{LWE}_{\cdot, N, Q}(\vec{s}_{\text{ext}}, \mathbf{m})$ with $\mathbf{m} = \mathbf{m}[1]$; that is, the LWE sample encodes the constant coefficient of the polynomial \mathbf{m}.</p>

Table 2: New parameter choices.

Set	Multi-Value	BR Key				KS Key		
		Q	N	ℓ_{brKey}	SD	n	ℓ_{ksKey}	SD
<code>tfhe-11-bin</code>	×	2^{48}	2^{11}	2	3.2	912	6	2^{26}
<code>tfhe-11-multi</code>	✓	2^{51}	2^{11}	3	3.2	950	6	2^{18}
<code>tfhe-12-multi</code>	✓	2^{50}	2^{12}	6	3.2	950	6	2^{18}

Dec(sk, ct). To decrypt an LWE sample $\vec{c} \in \mathbb{Z}_Q^{N+1}$, we compute $\langle \vec{c}, \vec{s}_{\text{ext}} \rangle = \frac{Q}{p} \cdot m' + e \in \mathbb{Z}_Q$, rescale and round the result obtaining

$$\left\lfloor \frac{p}{Q} \left(\frac{Q}{p} \cdot m' + e \right) \right\rfloor = m'$$

if $|e| \leq \frac{Q}{2p}$.

3.4 New Parameters for Multi-Value Bootstrapping

We choose our parameter sets to target 128-bit security for the LWE and RLWE samples. The parameters are listed in Table 2. We estimate the security using the latest commit of the Lattice Estimator [APS15]. We also include a Python script to estimate the statistical security. In Table 2, we specify three parameter sets. The `tfhe-11-bin` parameter set is based on previous work by Kluczniak [Klu22a] and chosen for binary ciphertexts. The parameter sets `tfhe-11-multi` and `tfhe-12-multi` are new parameter sets to support multi-value bootstrapping for 3-bit and 4-bit LUTs, respectively.

We choose our parameters according to the following strategy. For the bootstrapping key, we choose two rings, one with dimension $\text{deg} = 2^{11}$ and one with dimension $\text{deg} = 2^{12}$. The idea is to choose the highest modulus such that the RLWE problem remains 128-bit secure according to the Lattice Estimator [APS15] and the modulus is below 51-bits to allow for faster multiplication of ring elements using the HEXL library [BKS⁺21]. The larger ring gives us a larger group of the roots of unity, and we thus correctly process larger messages. For the LWE parameters, we set $n = 950$ for both rings and a uniform binary secret key since there are asymptotic reductions from binary LWE to LWE with uniform keys; we also check this parameter setting with the Lattice Estimator to ensure 128-bit security. We stress that we choose the secret key vector uniformly from the binary distributions. In particular, we do not use sparse secret keys and we do not fix the hamming weight, but there are algorithms to handle other key distributions [DM15, LMK⁺23]. However, these bootstrapping algorithms are usually slightly slower or require larger bootstrapping keys. Finally, we choose the decomposition bases to minimize the number of polynomial multiplications ℓ while at the same time preserving correctness with a probability of at most 2^{-80} for a faulty bootstrapping.

Based on Table 2, we can conclude that `tfhe-11-bin` will require the least amount of polynomial multiplications. In particular, recall that the number of polynomial multiplications is given by $n \cdot (2 \cdot (\ell_{\text{brKey}} + 1))$. For `tfhe-11-bin`, we need 5472 polynomial multiplications while for `tfhe-11-multi` and `tfhe-12-multi`, we need 7600 and 13300 polynomial multiplications, respectively. Furthermore, note that the ring dimension in `tfhe-12-multi` is doubled compared to the other parameter sets. Hence, a polynomial multiplication in this ring will be slower. We provide benchmarking results confirming these observations in Section 4.

3.5 Circuit Mapping

We now move from the bottom layer of our toolchain, the FHE scheme and its capabilities, to the post-synthesis layer above, that is to the netlist storing the synthesis output \mathcal{C}_2 and how the Boolean circuit is translated to a homomorphic circuit. Due to our efficient and secure multi-value bootstrapping, we upgrade homomorphic gate capabilities and are able to compute m outputs within a single bootstrapping via f_m . This enables a more generic approach to circuit mapping, which we describe in the following.

Recall that we assume a \mathcal{C}_2 which only consists of gates \mathcal{G}_n . To convert a Boolean circuit \mathcal{C}_2 to a homomorphic circuit \mathcal{C} , we compose the gate inputs via the affine part of the FHE scheme and translate each gate \mathcal{G}_n to a function f . Let $k = \max n$ be the maximum fan-in for all gates \mathcal{G}_n in \mathcal{C}_2 . We set $p = 2^k$ and compose the input to f as $x = \sum_{i=0}^{n-1} 2^i \cdot x_i$ homomorphically for encrypted values $x_i \in \mathbb{F}_2$. Theoretically, we could encrypt any value $x_i \in \mathbb{Z}_p$, but by only encrypting single bits, we have $x \in \mathbb{Z}_p$ and thus x serving as valid input to f by definition. Finally, we map the gate functionality of \mathcal{G}_n to a function f and compute the homomorphic equivalent as $f(x)$.

For example, consider the gate $\text{AND} : \mathbb{Z}_2^2 \rightarrow \mathbb{Z}_2, (I_0, I_1) \mapsto I_0 \cdot I_1$. We define f_{AND} mapping the encrypted plaintexts $\{0, 1, 2, 3\}$ to $\{0, 0, 0, 1\}$, respectively, and $f_{\text{AND}}(I_0 + 2I_1)$ computes an AND gate homomorphically. Note that, while the output domain of f is \mathbb{Z}_p , we will output only bits in $\{0, 1\}$. For a circuit \mathcal{C}_2 with gates $\mathcal{G}_{n,m}$, we trivially extend our mapping to a function f_m with m distinct functions f , one for every output bit.

3.6 Optimization: LUT Grouping

In the following, we introduce our first optimization performed during post-synthesis processing. Synthesis outputs \mathcal{C}_2 only consisting of single LUT gates $\mathcal{G}_{n,1}$, and our first optimization is constructing gates $\mathcal{G}_{n,m}$ to reduce the number of gates. As described in Subsection 3.5, we then trivially map these to homomorphic gates computing f_m via multi-value bootstrapping, thus reducing the number of bootstrappings compared to before. The straightforward case is two gates \mathcal{G}_n with exactly the same inputs in the same order which we can group as $\mathcal{G}_{n,2}$. There are, however, other opportunities enabling us to group different LUTs into a single LUT with multiple outputs. More specifically, we can group $\mathcal{G}_{n'}$ to \mathcal{G}_n as long as (1) the inputs of $\mathcal{G}_{n'}$ are a subset of the inputs of \mathcal{G}_n (or the other way around), and (2) grouping $\mathcal{G}_{n'}$ to \mathcal{G}_n does not introduce cycles to \mathcal{C}_2 .

Condition (1) follows relatively straightforward from our multi-value bootstrapping technique: We compute the blind rotation for an encrypted input $x \in \mathbb{Z}_p$ to the function f_m which, for a gate $\mathcal{G}_{n,m}$ and encrypted values $x_i \in \mathbb{Z}_2$, is composed as $x = \sum_{i=0}^{n-1} 2^i \cdot x_i$. We do have to be careful when grouping the individual f to a single f_m , however, as the x_i for all gates have to be in the same position i within a grouping. During post-synthesis optimization, we thus first extend each gate in the grouping to have all x_i as input, adding new inputs as required without modifying the gate output. Second, we order the inputs according to their unique IDs in HAL always composing the same x ; hence, we compute $\mathcal{G}_{n,m}$ correctly.

Condition (2) ensures that \mathcal{C}_2 still matches our definition from Subsection 2.2 preventing cycles in \mathcal{C}_2 as we otherwise could not evaluate it anymore. Consider the following scenario: Two inputs to some gates \mathcal{G}_2 and \mathcal{G}_3 are the same while the third input to \mathcal{G}_3 is the output of \mathcal{G}_2 ; here, grouping would introduce a cycle. This example generalizes to more complicated settings with additional gates in between and detecting such cases specifically is rather difficult; it would require checking all inputs to gates which are directly or indirectly connected to the output of a grouping candidate. We therefore choose a simpler, although non-optimal method, and only group $\mathcal{G}_{n'}$ to \mathcal{G}_n if \mathcal{G}_n appears before $\mathcal{G}_{n'}$ in the topological order and the inputs of $\mathcal{G}_{n'}$ are a subset of \mathcal{G}_n (but not the other way around).

An example for LUT grouping where both conditions are at play is depicted in Figure 2.

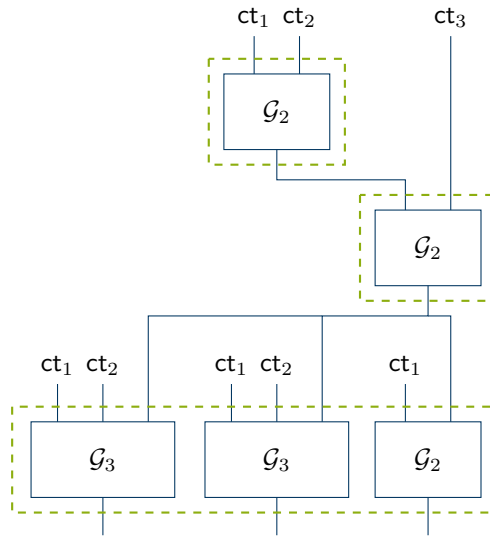


Figure 2: Grouping gates in an example circuit by matching the inputs while preserving topological order, each group is surrounded by a dashed line in green. The gate \mathcal{G}_2 in the top row cannot be grouped with the bottom row as it would introduce a cycle.

Here, the two gates \mathcal{G}_3 in the bottom row have the same inputs and hence can be grouped. Additionally, we can add \mathcal{G}_2 in the bottom row to the same group, as its inputs are a subset (we assume it appears later in the sorted graph). \mathcal{G}_2 in the middle row has a unique input combination and is its own group. As for the top row, the inputs of \mathcal{G}_2 are a subset of the inputs from the bottom row. But, adding it to the group at the bottom would introduce a cycle, and it thus remains its own group.

3.7 Optimization: Adder Substitution

Adder substitution, our second optimization, is a technique to improve arbitrary-width integer addition in Boolean-based schemes as for example as part of Concrete [Zam22], an FHE transpiler using instruction mapping. While the technique is known, there are—to the best of our knowledge—no existing descriptions in the scientific literature and we are the first to adapt this technique to the synthesis-based approach. It is based on the following observation: Using LUTs to realize additions is generally more costly than using the native addition capabilities of the FHE scheme (the same idea can be applied to transpilation based on instruction mapping). For example, with $p = 2^3$, a two-bit addition with LUT grouping requires two multi-value bootstrappings while adder substitution reduces it to one bootstrapping as depicted in Figure 3.

At its core, adder substitution takes place during synthesis moving us up another layer in the toolchain, but for simplicity, we split its implementation into two steps, one performed during synthesis in Yosys and one during post-synthesis processing in HAL: First, we hook into the synthesis process mapping additions and subtractions from the Yosys IR to custom adder gates. In the second step, we arrange these adders depending on the plaintext modulus p . In the following, we explore both steps in more detail.

During synthesis, we instruct Yosys to use single-bit full adder gates

$$\text{FA} : \mathbb{F}_2^3 \rightarrow \mathbb{F}_2^2, (x, y, c_i) \mapsto (s, c_o) = (x + y + c_i, x \cdot y + c_i \cdot (x + y))$$

for addition. Since an addition $x + y$ in the Yosys IR has arbitrary width, we start with a single-bit addition by adding the least significant bits x_0 and y_0 with $c_i = 0$, building up




Figure 3: Comparing a two-bit addition $s = x + y$ with a three-bit output based on LUTs for computation (left) and based on ciphertext addition with decomposition (right) with $p = 2^3$. The former requires two multi-value bootstrappings while the latter requires one.

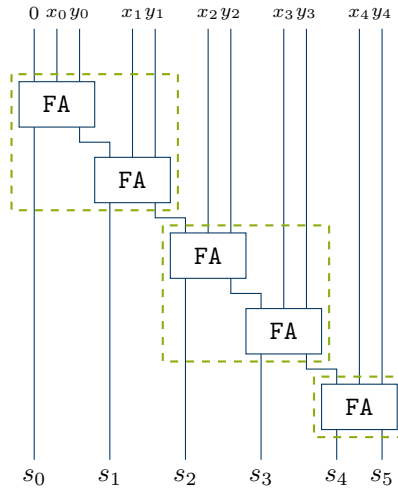


Figure 4: Arranging FA gates for a 5-bit addition $s = x + y$ with $p = 2^3$ and $c_i = 0$, each arrangement requiring a bootstrapping is marked by a dashed green box.

to an arbitrary-width addition by connecting the FA gates via the carry bits as required. For a subtraction $x - y$, we place inversion gates

$$\text{INV} : \mathbb{F}_2 \rightarrow \mathbb{F}_2, \quad I \mapsto O = I + 1$$

for y and set $c_i = 1$ for the least significant bit, this corresponds to a negation in the two's complement representation. Note that, for a ciphertext ct , computing inversion as $1 - \text{ct}$ does not require a bootstrapping and is considered free (see also Subsection 3.8).

During post-synthesis processing, we identify adder chains and form arrangements with $\log p - 1$ gates each. Using $\log p - 1$ gates ensures that addition does not wrap around modulo p , which we require for correctness. The FA for the least significant bit serves as root of a chain and is easy to identify with a constant $c_i \in \mathbb{F}_2$. Figure 4 contains an example for a 5-bit addition with $p = 2^3$.

As an additional optimization specific to the synthesis-based approach, we exclude constant addition or subtraction from adder substitution. The reason is rather straightforward: Since the individual bits of one operand are known at synthesis time, they can be folded with surrounding logic very efficiently and using our custom FA gate would result in worse performance. We evaluate the effect of constant bits on adder substitution in more detail in Section 4.

Another possible optimization is reserving more than one bit for carries. For example, with $\log p = 4$ bits and an addition of four values, we could add 2 bits from all four values instead of computing three individual additions using 3 bits from each value. We do not implement this because addition in the Yosys IR always has exactly two inputs. Note that this optimization can be applied to each addition in a circuit individually.

3.8 Additional Optimizations

Beyond the improvements achieved by LUT grouping and adder substitutions, we employ three further minor optimizations:

- We use the well-known fact that inversion of a ciphertext `ct` encrypting a value in \mathbb{F}_2 corresponds to the affine function $1 - ct$, also known as *free inversion* optimization; hence, no bootstrapping is required (this is also applicable to transpilation based on instruction mapping).
- We employ *LUT conversion*, during which we convert gates \mathcal{G}_1 to affine functions on a ciphertext without performing a bootstrapping. The mapping is trivial as a gate \mathcal{G}_1 either computes the identity function or an inversion.
- *Type sorting* is applied during the topological sorting of the circuit graph. We split the gates according to their depth and sort each layer with respect to gate types. Most importantly, we sort the LUTs according to the number of inputs in descending order. This ensures that we do not miss out on LUT groupings within a layer, as these cannot introduce cycles to the graph.

4 Evaluation

In the following, we evaluate our contributions with a multitude of examples; most designs stem from the FHE transpiler [GSPH⁺21], the remaining ones are our own addition. Overall, we use the following examples: Compute the sum of 3-bit, 4-bit or 32-bit integers, respectively (`add3`, `add4`, `add32`); calculate the addition, subtraction or multiplication of two 16-bit integers (`calc`); compute the sum of a 4-bit integer with a constant (`const4`); apply blurring to a small image (`img-blur`); apply a ricker wavelet transformation to a small image (`img-ricker`); apply sharpening to a small image (`img-sharp`); compute a rectified linear unit function (`relu`); compute the square root of a 16-bit integer (`sqr`); reverse an array of up to eight characters (`strrev`); compute the sum over a one-dimensional array of structs (`structs1d`); compute the sum over a three-dimensional array of structs (`structs3d`); compute the sum over a three-dimensional array of integers (`sum3d`).

For these examples, we target several architectures: The Yosys script from the FHE transpiler [GSPH⁺21] without optimizations (`prev`), our Yosys script with a maximum LUT-size of X bit without optimizations (`noX`), `noX` with adder substitution (`noXfa`), `noX` with LUT grouping including type sorting (`grpX`), `noX` with all optimizations (`optXfa`), and `optXfa` without adder substitution (`optX`).

4.1 Evaluating LUT Grouping

To evaluate LUT grouping, we compare the number of bootstrappings for architectures `noX` with architectures `grpX` for $X \in \{2, 3, 4\}$, respectively, in Table 3. For the chosen use cases, we reduce the number of bootstrappings by up to 75% and, for most use cases, by at least 35%. On average, using LUT grouping reduces the number of bootstrappings by almost 40%. Due to its effectiveness with only a slight overhead in the parameters, we recommend to always enable LUT grouping to improve performance for FHE circuit synthesis.

Table 3: Number of bootstrappings for all examples evaluating our two main optimizations using the architectures `noX`, `noXfa`, and `grpX` for $X \in \{2, 3, 4\}$.

	no2	no2fa	grp2	no3	no3fa	grp3	no4	no4fa	grp4
<code>add3</code>	12	3	7	6	2	3	5	1	2
<code>add4</code>	17	4	10	8	2	4	7	2	3
<code>add32</code>	163	32	102	63	16	32	74	11	34
<code>calc</code>	884	857	655	491	479	350	407	378	242
<code>const4</code>	5	5	2	4	4	1	4	4	1
<code>img-blur</code>	316	74	193	146	37	93	126	28	74
<code>img-ricker</code>	334	91	212	145	48	94	141	39	78
<code>img-sharp</code>	194	76	125	94	43	58	73	31	45
<code>relu</code>	16	16	15	15	15	15	15	15	15
<code>sqrt</code>	224	334	183	107	179	90	83	132	64
<code>strrev</code>	806	806	744	364	364	341	286	286	266
<code>structsid</code>	157	31	98	61	16	31	73	11	34
<code>structs3d</code>	1056	217	654	460	112	243	445	77	236
<code>sum3d</code>	926	203	553	396	105	199	381	74	192

```
[...]
assign sub_1935 = sel_1912 - ({1'h0, sel_1910, 13'h0000} | 16'h1000);
[...]
assign sub_1962 = sel_1939 - ({1'h0, sel_1937, 11'h000} | 16'h0400);
[...]
assign sub_1989 = sel_1966 - ({1'h0, sel_1964, 9'h000} | 16'h0100);
[...]
assign sub_2016 = sel_1993 - ({1'h0, sel_1991, 7'h00} | 16'h0040);
[...]
assign sub_2043 = sel_2020 - ({1'h0, sel_2018, 5'h00} | 16'h0010);
[...]
assign sub_2070 = sel_2047 - ({1'h0, sel_2045, 3'h0} | 16'h0004);
[...]
```

Listing 1: Subtractions in the `sqrt` example containing constant bits.

4.2 Evaluating Adder Substitution

For adder substitution, we compare the results for `noX` with `noXfa` for $X \in \{2, 3, 4\}$, respectively, in Table 3. Here, results are more mixed than before and we can make multiple interesting observations. As expected, for use cases without additions such as `strrev`, the number of bootstrappings stays the same. For use cases with lots of additions, however, improvements are much more drastic compared to before, and the number of bootstrappings is reduced by up to 85% and, on average, we reduce the number of bootstrappings by almost 45%.

Nevertheless, in the use case `sqrt`, we actually perform worse with our optimization. In Listing 1, we extract the culprits for this result. The subtractions performed depend on many constant bits for the second input. However, we cannot detect this during synthesis in Yosys as the built-in constant folding is not aggressive enough to mark the appropriate subset of input bits as constant. Since folding additions and subtractions with many constant input bits into LUTs is relatively efficient, the default LUT optimizations outperform adder substitution.

Overall, as efficient arithmetic is one of the main selling points of word-based schemes compared to Boolean-based schemes, we believe that the above optimization can be an important step to gain efficient arithmetic while keeping flexibility for FHE computations.

Table 4: Performance of the new parameter sets.

Set	Boot. [s]	brKey [MB]	ksKey [MB]
<code>tfhe-11-bin</code>	0.24	44.8	78.5
<code>tfhe-11-multi</code>	0.29	81.7	81.8
<code>tfhe-12-multi</code>	0.58	217.9	163.6

Table 5: Execution time in seconds including the corresponding rounded speed-ups compared to `prev` for all examples and a selection of synthesis processes with optimizations on our benchmarking setup.

	prev	opt2		opt3		opt3fa		opt4	
	t	t'	t/t'	t'	t/t'	t'	t/t'	t'	t/t'
<code>add3</code>	1.69	1.07	1.6	0.85	2.0	0.65	2.6	1.90	0.9
<code>add4</code>	2.32	1.50	1.5	1.13	2.1	0.74	3.1	2.73	0.8
<code>add32</code>	21.90	14.60	1.5	9.06	2.4	5.87	3.7	26.20	0.8
<code>calc</code>	106.00	77.10	1.4	71.30	1.5	72.30	1.5	123.00	0.9
<code>const4</code>	0.59	0.36	1.6	0.33	1.8	0.33	1.8	0.94	0.6
<code>img-blur</code>	40.10	25.50	1.6	20.90	1.9	10.50	3.8	46.60	0.9
<code>img-ricker</code>	43.00	27.80	1.5	21.80	2.0	12.40	3.5	48.50	0.9
<code>img-sharp</code>	23.10	16.30	1.4	13.00	1.8	8.68	2.7	27.60	0.8
<code>relu</code>	4.17	2.45	1.7	3.66	1.1	3.67	1.1	9.66	0.4
<code>sqrt</code>	24.00	21.10	1.1	18.60	1.3	34.60	0.7	33.50	0.7
<code>strrev</code>	82.70	88.70	0.9	70.60	1.2	70.80	1.2	137.00	0.6
<code>structs1d</code>	21.20	14.20	1.5	8.79	2.4	5.81	3.6	26.00	0.8
<code>structs3d</code>	138.00	86.80	1.6	59.00	2.3	32.90	4.2	152.00	0.9
<code>sum3d</code>	120.00	73.00	1.6	47.80	2.5	29.10	4.1	121.00	1.0

For now, we suggest transpiling circuits with and without adder substitutions and choosing the better performing option (this can happen automatically); we discuss possibilities for future work in Subsection 5.4 to avoid such scenarios.

4.3 Comparison with the FHE Transpiler

For comparison with the FHE transpiler, we use benchmarks with the actual parameters sets and FHE-Deck [fhe23] as library backend. We run benchmarks on a Windows Subsystem for Linux with kernel 5.15.146.1-microsoft-standard-WSL2 and an AMD Ryzen 9 7900X3D processor at 4.40 GHz featuring 12 cores. Our system has 16 GiB of available memory. For the new parameter sets, we summarize our results in Table 4 regarding runtime and memory consumption, confirming our observations from Subsection 3.4. For `prev`, we use the parameter sets `tfhe-11-bin` while we use `tfhe-11-multi` for `opt2`, `opt3`, and `opt3fa`. For `opt4`, we use `tfhe-12-multi`.

Our benchmarking results for all examples are summarized in Table 5. As expected, the execution time highly correlates with the number of bootstrappings. In most cases, we receive the best speed-ups for 3-bit LUTs. But, there are exceptions such as the `relu` example where 2-bit LUTs perform the best. Using 4-bit LUTs is generally not worth due to the increased polynomial degree and thus the longer bootstrapping time. For 3-bit LUTs, we outperform the FHE transpiler by up to $4.2\times$ with adder substitution and by up to $2.5\times$ without.

5 Discussion

In the following, we put our work in the context of current research, first discussing related work on FHEW-like implementations followed by related work on FHE circuit synthesis. Afterward, we discuss limitations of our optimizations and explore multiple opportunities for future work to further optimize the tool-based generation of circuits for FHE.

5.1 Related Work on FHEW-like Implementations

In Table 6, we roughly compare different libraries implementing FHEW-like schemes. A distinguishing feature of FHE-Deck is the support for correct and secure parameter sets for multi-value bootstrapping. Both FHE-Deck and TFHE-rs support different algorithms for functional bootstrapping (also known as programmable bootstrapping). In particular, FHE-Deck supports the full domain bootstrapping algorithm based on work by Liu, Micciancio, and Polyakov [LMP22] while THFE-rs supports the algorithm by Chilotti et al. [CLOT21]. Moreover, both libraries support simple padding-based functional bootstrapping. The TFHE library [CGGI16b] supports only binary gates. Open-FHE [ABBB⁺22], which is derived from PALISADE [PAL21], implements binary as well as full domain bootstrapping based on work by Liu, Micciancio, and Polyakov [LMP22].

There are implementations for LUT evaluation [CGGI17]. However, the techniques are vastly different, as the authors evaluate a LUT on so-called RGSW ciphertexts requiring numerous bootstrapping invocations (the number depends on the chosen parameters) for each output bit of the LUT. In contrast, we focus on computing LUTs using a single bootstrapping invocation. Finally, the amortized bootstrapping technique by Micciancio and Sorrel [MS18] and its improvement [GPvL23] compute functional bootstrapping over many input ciphertexts at a cheaper cost than bootstrapping ciphertexts separately. In particular, the functional difference is that we amortize computation for many output functions on the same input ciphertexts while they compute the same functions on multiple ciphertexts. Combining both amortization techniques in a practical way is an interesting open problem for future work.

Multi-Value Bootstrapping

For multi-value bootstrapping, Carpov, Izabachène, and Mollimard [CIM19] factorize a polynomial used as input to the blind rotation, returning factors \mathbf{w}_0 and \mathbf{w}_1 where \mathbf{w}_1 has a low norm. Aligned with our notation, they would set \mathbf{w} to \mathbf{w}_0 and insert \mathbf{w}_1 into the vector $\vec{\mathbf{v}}$. Essentially, the factorization works as follows: They set $\mathbf{w}_0 = \sum_{i=1}^N X^i$ and $\mathbf{w}_1 = \sum_{i=1}^N t'_i X^i$. Then, to compute the t'_i coefficients, they build and solve a large system of N linear equations. Commonly, the polynomial degree is a power-of-two $N \geq 2^{11}$ and Gaussian elimination runs in cubic time in the number of variables. Solving such a system (if solvable) may take considerable time. In contrast, our method only requires N invocations of a function $f^{(i)}$ that we want to encode in a polynomial in $\vec{\mathbf{v}}$.

The TFHE library implements multi-value bootstrapping in a separate branch. The authors only test the bootstrapping for randomly chosen rotation polynomials \mathbf{w}_0 and \mathbf{w}_1 . Unfortunately, we couldn't find the implementation of the linear system solver. Nevertheless, for $N = 2^{14}$, as suggested by [CIM19], we can roughly calculate that the number of modular multiplications in the Gaussian elimination algorithm⁷ as $> 2^{41.415}$, a considerable time in practice. We also need to assume that the system is solvable which requires elements of the matrix that we build for Gaussian elimination to be invertible modulo Q , which, with high probability, will not be the case if Q is a power-of-two as in many implementations (this includes their own implementation; note that we may choose

⁷Recall that Gaussian elimination requires $N(N+1)/2$ divisions, $(2N^3 + 3N^2 - 5N)/6$ multiplications, and $(2N^3 + 3N^2 - 5N)/6$ subtractions modulo Q .

Table 6: Functionality comparison of different FHE libraries for FHEW-like schemes. For functional bootstrapping, we denote as \circ the plain FHEW/TFHE algorithm to evaluate boolean gates. By \bullet , we denote a full domain functional bootstrapping algorithm [YXS⁺21, CLOT21, LMP22, KS22]. By \bullet , we denote support for our improved functional bootstrapping algorithm. In this comparison, a high-level interface for functional bootstrapping is required.

Library	Language	Bootstrapping	
		Functional	Multi-Value
FHE-Deck	C++	\bullet	\checkmark
Open-FHE	C++	\bullet	\times
TFHE	C++	\circ	\times
tfhe-rs/CONCRETE	Rust	\bullet	\times

Table 7: Comparison of methods to choose \mathfrak{w} and $\vec{\mathfrak{v}}$.

Method	Time	Modulus
[CIM19]	$O(N^3)$	Prime
Ours	$O(N)$	Any

a prime modulus to fix this which is not a limitation in our approach). We compare the time complexity and the restrictions on the modulus in Table 7.

5.2 Related Work on Circuit Synthesis

There are a couple of previous works on circuit synthesis for FHEW-like schemes, the previously mentioned FHE transpiler [GSPH⁺21] as well as two earlier works named Cingulata, originally released under the name Armadillo [CDS15], and the E3 framework [CMG⁺18]. In Cingulata, the authors divide their toolchain in three parts: the front-end translating C++ code to a Boolean circuit, the middle-end optimizing the circuit and the back-end transpiling the circuit to an FHE library. The mapping from C++ to a Boolean circuit in the front-end defers optimizations to the middle-end based on ABC [Mis], which we also use as part of our toolchain via Yosys. No other optimizations are performed. The E3 framework also uses hardware tooling for transpilation, however, no details regarding optimizations are available in their publication.

The FHE transpiler currently improves upon all previously known work and thus serves as a good foundation to evaluate new research ideas. Common compiler optimizations such as constant folding or dead code elimination are performed in the XLS-based high-level synthesis layer. However, to the best of our knowledge, the only FHE-specific optimization currently performed is rather trivial treating inversion as free for Boolean-based circuits. For LUT-based circuits, there is currently no post-processing implementing this optimization.

After the initial release of our work, Guan et al. [GMZ⁺24] published the framework AutoHoG which also uses multi-value bootstrapping for homomorphic gate computations. Unfortunately, we were unable to replicate their results for a fair comparison since we could not derive certain parameters or implementation details from the description in their paper. For example, they do not specify the key and error distributions and we did not manage to instantiate TFHE securely⁸ with the remaining parameters as specified in their work. Singh et al. [SDC⁺23] create homomorphic three-bit AND, XOR, and NAND gates generalizing the well-known two-bit method from [DM15]. Here, we were also unable to derive all required parameters and implementation details for a fair comparison.

⁸using the Lattice Estimator [APS15]

Importantly, our LUT-based approach with multi-value bootstrapping encapsulates their approach because it enables generic gates $\mathcal{G}_{n,m}$ with a single bootstrapping.

5.3 Limitations

There are some limitations for our proposed optimizations. First, as highlighted in Subsection 4.2, using adder substitution can result in worse performance when the inputs contain many constant bits which can be non-detectable using our current approach. Therefore, a user has to manually check for the better circuit. One solution and useful contribution in future work would be improving constant bit detection and constant folding in the Yosys IR. Second, although using three-bit ciphertexts tends to provide the largest speed-ups, sometimes other bit sizes can be more beneficial such as using two-bit ciphertexts for the `relu` example. Exploring the root causes and detecting such cases, especially if also done for subcircuits, would further improve circuit synthesis for FHE.

5.4 Future Work

In our opinion, an important observation is that the current state-of-the-art in bootstrapping for FHEW-like schemes can still be improved upon. We also believe that there is still room for improvement regarding performance for TFHE implementations as well as regarding usability for currently available libraries, including, but not limited to, FHE-Deck. As for circuit synthesis, our work is a first step in FHE-specific optimizations and we believe that there is a multitude of other possibilities making automatically generated circuits more efficient. For example, multi-value bootstrapping greatly benefits from LUTs with the same inputs which current hardware tooling is not optimizing for. Overall, we are looking forward to future work in the area of circuit synthesis for Boolean-based FHE schemes.

6 Conclusion

In this work, we improve performance and usability of FHEW-like schemes by extending the current state-of-the-art in bootstrapping as well as circuit synthesis. To improve performance, we significantly simplify the bootstrapping idea proposed by Carpov, Izabachène, and Mollimard [CIM19] and provide the first efficient implementation for multi-value bootstrapping. With respect to usability, we provide new and secure parameter sets for multi-bit encryptions, which can be used by researchers and developers alike, and implement a high-level interface for multi-value bootstrapping in the open-source library FHE-Deck [fhe23].

We provide a generalized model for mapping Boolean circuits to homomorphic circuits and introduce the first non-trivial FHE-specific optimizations for generating circuits from high-level code: LUT grouping and adder substitution. Using LUT grouping, generated circuits require almost 40% less bootstrappings on average and adder substitution reduces the number of required bootstrappings by up to 85%. Overall, using all optimizations results in up to $4.2\times$ faster execution times compared to previous synthesized circuits with state-of-the-art methods.

Acknowledgements

Johannes Mono and Tim Güneysu are supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2092 CASA - 390781972. Kamil Klucznik is supported by the German Ministry for Education and Research through funding for the project CISP-Stanford Center for Cybersecurity (Funding number: 16KIS0927).

References

- [ABBB⁺22] Ahmad Al Badawi, Jack Bates, Flavio Bergamaschi, David Bruce Cousins, Saroja Erabelli, Nicholas Genise, Shai Halevi, Hamish Hunt, Andrey Kim, Yongwoo Lee, Zeyu Liu, Daniele Micciancio, Ian Quah, Yuriy Polyakov, Saraswathy R.V., Kurt Rohloff, Jonathan Saylor, Dmitriy Suponitsky, Matthew Triplett, Vinod Vaikuntanathan, and Vincent Zucca. Openfhe: Open-source fully homomorphic encryption library. In *Proceedings of the 10th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, WAHC'22, pages 53–63, New York, NY, USA, 2022. Association for Computing Machinery.
- [AP14] Jacob Alperin-Sheriff and Chris Peikert. Faster bootstrapping with polynomial error. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 297–314, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.
- [BDF18] Guillaume Bonnoron, Léo Ducas, and Max Fillinger. Large FHE gates from tensored homomorphic accumulator. In Antoine Joux, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 18: 10th International Conference on Cryptology in Africa*, volume 10831 of *Lecture Notes in Computer Science*, pages 217–251, Marrakesh, Morocco, May 7–9, 2018. Springer, Heidelberg, Germany.
- [BGV14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. *ACM Trans. Comput. Theory*, 6(3):13:1–13:36, 2014.
- [BKS⁺21] Fabian Boemer, Sejun Kim, Gelila Seifu, Fillipe DM de Souza, Vinodh Gopal, et al. Intel HEXL (release 1.2). <https://github.com/intel/hexl>, September 2021.
- [Bra12] Zvika Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 868–886, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.
- [CDS15] Sergiu Carpov, Paul Dubrulle, and Renaud Sirdey. Armadillo: A compilation chain for privacy preserving applications. In *Proceedings of the 3rd International Workshop on Security in Cloud Computing*, SCC '15, page 13–19, New York, NY, USA, 2015. Association for Computing Machinery.
- [CGGI16a] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *Advances in Cryptology – ASIACRYPT 2016, Part I*, volume 10031 of *Lecture Notes in Computer Science*, pages 3–33, Hanoi, Vietnam, December 4–8, 2016. Springer, Heidelberg, Germany.

- [CGGI16b] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast fully homomorphic encryption library, August 2016. <https://tfhe.github.io/tfhe/>.
- [CGGI17] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. Faster packed homomorphic operations and efficient circuit bootstrapping for TFHE. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017, Part I*, volume 10624 of *Lecture Notes in Computer Science*, pages 377–408, Hong Kong, China, December 3–7, 2017. Springer, Heidelberg, Germany.
- [CIM19] Sergiu Carpov, Malika Izabachène, and Victor Mollimard. New techniques for multi-value input homomorphic evaluation and applications. In Mitsuru Matsui, editor, *Topics in Cryptology – CT-RSA 2019*, volume 11405 of *Lecture Notes in Computer Science*, pages 106–126, San Francisco, CA, USA, March 4–8, 2019. Springer, Heidelberg, Germany.
- [CKKS17] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. Homomorphic encryption for arithmetic of approximate numbers. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017, Part I*, volume 10624 of *Lecture Notes in Computer Science*, pages 409–437, Hong Kong, China, December 3–7, 2017. Springer, Heidelberg, Germany.
- [CLOT21] Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Improved programmable bootstrapping with larger precision and efficient arithmetic circuits for TFHE. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part III*, volume 13092 of *Lecture Notes in Computer Science*, pages 670–699, Singapore, December 6–10, 2021. Springer, Heidelberg, Germany.
- [CMG⁺18] Eduardo Chielle, Oleg Mazonka, Homer Gamil, Nektarios Georgios Tsoutsos, and Michail Maniatakos. E3: A framework for compiling c++ programs with encrypted operands. Cryptology ePrint Archive, Paper 2018/1013, 2018. <https://eprint.iacr.org/2018/1013>.
- [DM15] Léo Ducas and Daniele Micciancio. FHEW: Bootstrapping homomorphic encryption in less than a second. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 617–640, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- [Emb19] Embedded Security Group. HAL - The Hardware Analyzer. <https://github.com/emsec/hal>, 2019.
- [fhe23] Fhe-deck. <https://github.com/FHE-Deck>, September 2023.
- [FV12] Junfeng Fan and Frederik Vercauteren. Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144, 2012. <https://eprint.iacr.org/2012/144>.
- [GBA21] Antonio Guimarães, Edson Borin, and Diego F. Aranha. Revisiting the functional bootstrap in tfhe. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2021(2):229–253, Feb. 2021.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st Annual ACM Symposium on Theory of Computing*, pages 169–178, Bethesda, MD, USA, May 31 – June 2, 2009. ACM Press.

- [GMZ⁺24] Zhenyu Guan, Ran Mao, Qianyun Zhang, Zhou Zhang, Zian Zhao, and Song Bian. Autohog: Automating homomorphic gate design for large-scale logic circuit evaluation. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pages 1–1, 2024.
- [GPvL23] Antonio Guimarães, Hilder V. L. Pereira, and Barry van Leeuwen. Amortized bootstrapping revisited: Simpler, asymptotically-faster, implemented. Cryptology ePrint Archive, Report 2023/014, 2023. <https://eprint.iacr.org/2023/014>.
- [GSPH⁺21] Shruthi Gorantala, Rob Springer, Sean Purser-Haskell, William Lam, Royce Wilson, Asra Ali, Eric P. Astor, Itai Zukerman, Sam Ruth, Christoph Dibak, Phillipp Schoppmann, Sasha Kulankhina, Alain Forget, David Marn, Cameron Tew, Rafael Misoczki, Bernat Guillen, Xinyu Ye, Dennis Kraft, Damien Desfontaines, Aishe Krishnamurthy, Miguel Guevara, Irippuge Milinda Perera, Yurii Sushko, and Bryant Gipson. A general purpose transpiler for fully homomorphic encryption. Cryptology ePrint Archive, Report 2021/811, 2021. <https://eprint.iacr.org/2021/811>.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 75–92, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [Klu22a] Kamil Kluczniak. NTRU- ν -um: Secure fully homomorphic encryption from NTRU with small modulus. Cryptology ePrint Archive, Report 2022/089, 2022. <https://eprint.iacr.org/2022/089>.
- [Klu22b] Kamil Kluczniak. NTRU- ν -um: Secure fully homomorphic encryption from NTRU with small modulus. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 1783–1797, Los Angeles, CA, USA, November 7–11, 2022. ACM Press.
- [KS22] Kamil Kluczniak and Leonard Schild. Fdfb: Full domain functional bootstrapping towards practical fully homomorphic encryption. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023(1):501–537, Nov. 2022.
- [LMK⁺23] Yongwoo Lee, Daniele Micciancio, Andrey Kim, Rakyong Choi, Maxim Deryabin, Jieun Eom, and Donghoon Yoo. Efficient FHEW bootstrapping with small evaluation keys, and applications to threshold homomorphic encryption. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part III*, volume 14006 of *Lecture Notes in Computer Science*, pages 227–256, Lyon, France, April 23–27, 2023. Springer, Heidelberg, Germany.
- [LMP22] Zeyu Liu, Daniele Micciancio, and Yuriy Polyakov. Large-precision homomorphic sign evaluation using FHEW/TFHE bootstrapping. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part II*, volume 13792 of *Lecture Notes in Computer Science*, pages 130–160, Taipei, Taiwan, December 5–9, 2022. Springer, Heidelberg, Germany.
- [Mis] Alan Mishchenko. System for sequential logic synthesis and formal verification. <https://github.com/berkeley-abc/abc>.

- [MS18] Daniele Micciancio and Jessica Sorrell. Ring packing and amortized FHEW bootstrapping. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *ICALP 2018: 45th International Colloquium on Automata, Languages and Programming*, volume 107 of *LIPICs*, pages 100:1–100:14, Prague, Czech Republic, July 9–13, 2018. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.
- [PAL21] PALISADE Lattice Cryptography Library (release 1.11.5). <https://palisade-crypto.org/>, September 2021.
- [RAD78] Ronald Rivest, Len Adleman, and Michael Dertouzos. On data banks and privacy homomorphism, 1978.
- [SDC⁺23] Animesh Singh, Smita Das, Anirban Chakraborty, Rajat Sadhukhan, Ayantika Chatterjee, and Debdeep Mukhopadhyay. FHEDA: efficient circuit synthesis with reduced bootstrapping for Torus FHE. *IACR Cryptol. ePrint Arch.*, page 1310, 2023.
- [YXS⁺21] Zhaomin Yang, Xiang Xie, Huajie Shen, Shiyong Chen, and Jun Zhou. TOTA: Fully homomorphic encryption with smaller parameters and stronger security. Cryptology ePrint Archive, Report 2021/1347, 2021. <https://eprint.iacr.org/2021/1347>.
- [Zam22] Zama. Concrete: TFHE Compiler that converts python programs into FHE equivalent, 2022. <https://github.com/zama-ai/concrete>.