

BASALISC: Flexible Asynchronous Hardware Accelerator for Fully Homomorphic Encryption

Robin Geelen*
robin.geelen1@kuleuven.be
imec-COSIC KU Leuven
Leuven, Belgium

Brian Huffman
huffman@galois.com
Galois, Inc.
Portland, OR, USA

Daniel Wagner
dmwit@galois.com
Galois, Inc.
Portland, OR, USA

Michiel Van Beirendonck*
michiel.vanbeirendonck@kuleuven.be
imec-COSIC KU Leuven
Leuven, Belgium

Tynan McAuley
tynan@niobiummicrosystems.com
Niobium Microsystems
Portland, OR, USA

Georgios Dimou
georgios@niobiummicrosystems.com
Niobium Microsystems
Portland, OR, USA

Hilder V. L. Pereira
hildervitor.limapereira@kuleuven.be
imec-COSIC KU Leuven
Leuven, Belgium

Ben Selfridge
benselfridge@galois.com
Galois, Inc.
Portland, OR, USA

Ingrid Verbauwhede
ingrid.verbauwhede@kuleuven.be
imec-COSIC KU Leuven
Leuven, Belgium

Frederik Vercauteren
frederik.vercauteren@kuleuven.be
imec-COSIC KU Leuven
Leuven, Belgium

David W. Archer
dwa@galois.com
Galois, Inc.
Portland, OR, USA

ABSTRACT

Fully Homomorphic Encryption (FHE) allows for secure computation on encrypted data. It enables a variety of theoretical and practical applications, but is still several orders of magnitudes too slow to be practical. We present BASALISC, an architecture family of FHE hardware accelerators that aims to substantially accelerate FHE computations in the cloud. BASALISC implements *Brakerski, Gentry, and Vaikuntanathan's* (BGV) scheme and supports a range of parameter sets. In contrast to many prior studies, we directly support and implement BGV bootstrapping – the noise removal capability necessary to support arbitrary-depth computation.

BASALISC exploits data representation in residue number systems and number-theoretic transforms to realize massive FHE parallelism. We propose a new generalized version of bootstrapping that can be implemented with optimized Montgomery multipliers that cost 46% less in silicon area and 40% less in power consumption versus traditional approaches. BASALISC is a *Reduced Instruction Set Computing* (RISC) architecture with a four-layer memory hierarchy, including a two-dimensional conflict-free inner memory layer that enables 32 Tb/s radix-256 number-theoretic transform (NTT) computations without pipeline stalls. Our conflict-resolution data permutation hardware is re-used to compute BGV automorphisms without additional hardware and without throughput penalty. BASALISC additionally includes a custom multiply-accumulate unit familiar in Digital Signal Processing (DSP) architectures, with which we accelerate tight BGV key switching loops. The BASALISC computation units and inner memory layers are designed in asynchronous logic, allowing them to run at different speeds to optimize each function. BASALISC is designed for *Application-Specific Integrated Circuit* (ASIC) implementation with

a 1 GHz operational frequency, and is already underway toward tape-out with a 150mm² die size in a 12nm Global Foundries process.

The BASALISC toolchain comprises both a custom compiler and a joint performance and correctness simulator. We evaluate BASALISC in multiple ways: we study its physical realizability; we emulate and formally verify its core functional units; and we study its performance on a single iteration of logistic regression training over encrypted data. For this application, comprising from up to 900K high-level BASALISC instructions (including 513 bootstraps) down to 27B low-level instructions, we show a speedup of at least 2,025× over HELib - a popular software FHE library - running on a Xeon-class processor.

KEYWORDS

FHE, BGV, Hardware Accelerator, ASIC

1 MOTIVATION

Fully Homomorphic Encryption (FHE) [3, 18, 39] offers the promise of confidentiality-preserving computation over sensitive data in a variety of theoretical and practical applications, ranging from new cryptographic primitives to machine learning as a service. Unfortunately, the utility of FHE is severely limited by twin challenges of inefficient memory use and high computational overhead. The typical result - computation that runs many orders of magnitude slower than insecure computation - prevents broad adoption. Although new schemes have markedly improved FHE performance [6, 7, 11, 12], and highly optimized FHE libraries [1, 13, 22, 33, 43] are now available, FHE still remains orders of magnitude beyond acceptable performance limits for most potential applications.

*R. Geelen and M. Van Beirendonck contributed equally to this research.

In other computational domains where performance on general purpose processors is problematic, innovation has turned to purpose-built *accelerators*, tuned to exploit domain-specific characteristics of computation. DSP accelerators, arguably starting with the Texas Instruments TMS320 DSP family [28] in 1983, are perhaps the first example of this approach. More recently, Graphics Processing Units (GPUs) have become popular for accelerating video stream processing and hash function computation. Our FHE accelerator, BASALISC, follows this approach in pursuit of bringing the throughput of FHE computation within an order of magnitude relative to cleartext computation. In doing so, BASALISC employs many of the computational principles from both DSP and GPU accelerators.

We summarize the key contributions of BASALISC as follows:

- BASALISC accelerates BGV arithmetic for a range of parameters. BASALISC is a comprehensive RISC-like architecture with a three-level instruction set architecture (ISA) that allows for reasoning at diverse levels of executive abstraction. In contrast to prior architectures, BASALISC supports and implements bootstrapping to enable unlimited-depth FHE computations.
- We propose a novel version of bootstrapping that is compatible with Montgomery-friendly primes. In contrast to prior work, BASALISC instantiates its multipliers exclusively to these Montgomery-friendly primes, which saves 46% logic area and 40% power consumption.
- BASALISC implements a massively parallel radix-256 NTT architecture, using a conflict-free layout, a corresponding layout permutation unit, and a twiddle factor generator. These units are deeply interleaved with the on-chip memory and provide a total 32 Tb/s NTT throughput. In addition, we show that we can efficiently generalize the required layout permutation unit to compute BGV automorphisms *without additional silicon area*.
- BASALISC adopts a four-level memory hierarchy purpose-built to address common FHE memory bottlenecks, including a mid-level 64 MB on-chip ciphertext buffer (CTB). At the lowest level, a massively parallel multiply-accumulate unit with integrated 16-entry register file allows accelerating tight BGV key switching loops, asynchronously and independently of the CTB.
- BASALISC is placed and routed with 150mm² die size and 1 GHz operational frequency in a 12nm low-power Global Foundries process. Critical hardware logic is emulated and formally verified for correctness. We evaluate BASALISC on a bootstrapping benchmark and a logistic regression application, showing respectively 4,000× and 2,025× speedup over an HElib software reference.

2 PRELIMINARIES

2.1 Fully Homomorphic Encryption

Fully Homomorphic Encryption (FHE) provides a simple use model to securely outsource computation on sensitive data to a third party. Informally, the FHE model enables a user to encrypt its data m into a ciphertext $c = \text{Enc}(m)$, then send it to a third party, who can compute on c . The third party produces another ciphertext c'




Figure 1: FHE used in a typical commercial application.

encrypting $f(m)$ for some desired function f . We say that f was computed homomorphically.

In FHE, the third party receives only ciphertexts and the public key, but never the secret key that allows decryption. As a result, the sensitive inputs are protected under the security of the encryption scheme. Because the result of the computation remains encrypted, the output also remains unknown to the third party: only the holder of the secret key can decrypt and access it. This scenario is illustrated in Figure 1: the client generates a key pair, and shares the public key with the cloud server. Then the client sends an encryption of its data, which is processed homomorphically by the server, and finally sent back to the client.

To achieve security, the ciphertexts of all FHE schemes are noisy: during encryption, a small noise term is added to the input data. Decryption can still recover the correct result, provided that the noise is small enough. To evaluate a function homomorphically, we represent the function in terms of the operations provided by the scheme, typically addition and multiplication, and compute these operations on the encrypted inputs. Each operation increases the noise in the resulting ciphertext, so we can compute only a limited number of homomorphic operations before we reach the limit of decryption failure.

Because multiplications increase ciphertext noise much more than additions, we usually model noise growth by the number of sequential multiplications only. If we compute the product $\prod_{i=1}^L m_i$ homomorphically, then we say that the computation requires multiplicative depth $\lceil \log_2(L) \rceil$. This is accomplished by writing the product in a tree structure, with each leaf node representing one of the factors. In general, there is a trade-off between computational cost and tolerating a larger L : we can increase the parameters used to instantiate the scheme so that we obtain more multiplicative depth, but in doing so we make the homomorphic operations slower and the size of ciphertexts larger.

To support the computation of functions regardless of their multiplicative depth, FHE uses *bootstrapping*. This is an operation that reduces ciphertext noise by decrypting it homomorphically. Unfortunately, bootstrapping is very expensive, so its use is often minimized in practical applications. There are several techniques in the FHE literature to slow down the noise growth, and thus delay bootstrapping. In this work, we employ *key switching* and *modulus switching* [7]. We note that bootstrapping and key switching tend to heavily dominate computation and data movement costs of an application: in a simple 1,024-point, 10-feature logistic regression,

we see that these tasks account for over 95% of the computational effort and the vast majority of data movement.

The key challenges in designing an efficient FHE scheme are the high complexity of computation, the large ciphertext expansion factor (large polynomials with integer coefficients of 1000 bits or more), and the proportion of effort needed in bootstrapping (or delaying it) in sufficiently complex programs. In the remainder of this paper, we examine the magnitude of these challenges and how they impact our FHE accelerator.

2.2 The BGV Cryptosystem

BASALISC targets the homomorphic encryption scheme known as BGV [7]. Plaintexts and ciphertexts are represented by elements in the ring $\mathcal{R} = \mathbb{Z}[X]/(X^N + 1)$ with N a power of 2. Those elements are thus polynomials reduced modulo $X^N + 1$, and this modular reduction is implicit in our notation. BGV guarantees finite data structures by also reducing the coefficients: the plaintext space is computed modulo t (denoted \mathcal{R}_t), and the ciphertext space is a pair of elements modulo q (denoted \mathcal{R}_q^2). Reduction modulo m (with $m = t$ or q) is explicitly denoted by $[\cdot]_m$. It is always done symmetrically around 0, i.e., in the set $[-m/2, m/2) \cap \mathbb{Z}$.

As with traditional ciphers, BGV has encryption and decryption procedures to move between the *plaintext space* and the *ciphertext space*. These operations are never executed by the server doing out-sourced computation and therefore not implemented by BASALISC. However, it is necessary to explain the ciphertext format in order to understand homomorphic operations. A BGV ciphertext $(c_0, c_1) \in \mathcal{R}_q^2$ is said to encrypt plaintext $m \in \mathcal{R}_t$ under secret key s (which has small coefficients) if

$$c_0 + c_1 \cdot s = m + te \pmod{q} \quad (1)$$

for some element e that also has small coefficients. The term e is called the *noise*, and it determines if decryption returns the correct plaintext: as long as e has coefficients roughly smaller than $q/2t$, the expression $m + te$ does not overflow modulo q . We can therefore recover the plaintext uniquely as $m = [(c_0 + c_1 \cdot s)_q]_t$.

2.2.1 Basic Homomorphic Operations. Smart and Vercauteren observed that for $t = p^r$ with p an odd prime, the plaintext space \mathcal{R}_t is equivalent to \mathbb{Z}_t^ℓ for some ℓ that divides N [47]. This technique is referred to as *packing*, and it allows us to encode ℓ numbers into one plaintext simultaneously. Addition and multiplication over tuples in \mathbb{Z}_t^ℓ are then performed component-wise. As a result, one ciphertext can encrypt and operate on an entire tuple, which leads to significant performance gains and memory reductions in practice.

When BGV is used in conjunction with packing, we can define three basic homomorphic operations. Let (c_0, c_1) and (c'_0, c'_1) be two ciphertexts encrypting the tuples (m_1, \dots, m_ℓ) and (m'_1, \dots, m'_ℓ) .

- **Addition:** we compute $([c_0 + c'_0]_q, [c_1 + c'_1]_q)$. The encrypted plaintext is $(m_1 + m'_1, \dots, m_\ell + m'_\ell)$.
- **Multiplication:** we compute $([c_0 \cdot c'_0]_q, [c_0 \cdot c'_1 + c_1 \cdot c'_0]_q, [c_1 \cdot c'_1]_q)$. The resulting ciphertext is a vector of three elements, but this can be reduced back to two with a post-processing step called *key switching* (see later). The encrypted plaintext is $(m_1 \cdot m'_1, \dots, m_\ell \cdot m'_\ell)$.
- **Permutation:** we compute $(\phi_k(c_0), \phi_k(c_1))$, where the map ϕ_k is called an *automorphism*. It is parameterized by an

odd integer k , and defined as $\phi_k: c(X) \mapsto c(X^k)$. Gentry et al. [20] show that these automorphisms induce a permutation on the elements of the encoded tuple, so the output encrypts some permutation of (m_1, \dots, m_ℓ) . Although the resulting ciphertext only has two elements, we still need post-processing by means of key switching.

The validity of these three operations can simply be verified by observing their effect on Equation 1. We refer to Zucca [51] for a more detailed analysis, including noise growth of each operation.

2.2.2 Auxiliary Homomorphic Operations. Basic homomorphic operations lead to ciphertext expansion and noise growth. Take for example a product ciphertext: it consists of three elements and it is encrypted under (s, s^2) instead of s . The same problem occurs during permutation: the automorphism ϕ_k has a side effect on the secret key, so the resulting ciphertext is encrypted under $\phi_k(s)$. Also noise growth is an issue: the noise term in a product ciphertext, for example, has increased to $te \cdot e'$.

To prevent ciphertext expansion, switch between keys and slow down noise growth, BGV defines two auxiliary procedures:

- **Modulus switching:** given a ciphertext $(c_0, c_1) \in \mathcal{R}_q^2$ and a new modulus q' , we compute a ciphertext $(c'_0, c'_1) \in \mathcal{R}_{q'}^2$ that decrypts with respect to q' . Modulus switching also scales the noise by a factor of q'/q .
- **Key switching:** given a key switching matrix (\vec{k}_0, \vec{k}_1) and either a product ciphertext $(c_0, c_1, c_2) \in \mathcal{R}_q^3$ or a permuted ciphertext $(c_0, c_1) \in \mathcal{R}_q^2$, we compute a ciphertext $(c'_0, c'_1) \in \mathcal{R}_q^2$ that decrypts under Equation 1. Thus key switching brings the ciphertext back to its original format.

In summary, modulus switching is run before each multiplication to reduce the noise to its minimum level. Key switching is run after each permutation or multiplication to keep the ciphertext format consistent. Again we refer to Zucca [51] for a more detailed analysis.

2.2.3 Bootstrapping. When the entire noise budget of a ciphertext is consumed (equivalently, when the modulus q is depleted to its minimum value by successive modulus switchings), further homomorphic operations are no longer immediately possible. We can overcome this problem by means of a *bootstrapping* procedure that reduces the noise back to a lower level [18]. Bootstrapping “refreshes” a ciphertext by running decryption *homomorphically*: we first evaluate an adapted version of Equation 1, followed by coefficient-wise rounding. The currently most efficient bootstrapping technique for BGV is implemented in the HELib library [23].

2.2.4 Supported Parameter Sets. As a nod to Amdahl’s Law (“make the common case fast”), hardware optimization gains throughput benefits by supporting only a limited range of commonly used parameters. We start with the realization that at least 128-bit security must be supported if BASALISC is to be interesting to real-world users. Based on this observation, we choose a range of parameters that allows for an efficient implementation, while still retaining sufficient freedom for application design.

Some potential customers have indicated a desire for 256-bit security.

Table 1: BASALISC parameter ranges and examples.

Parameter	Range	Example
Security parameter	N/A	128 bits
Ring dimension N	512 – 65536	65536
Plaintext modulus p^r	≥ 2	127^3
Ciphertext packing ℓ	1 – 65536	64 slots
Max $\log_2(QP)$ for key switching	20 – 1782	1782 bits
Max $\log_2(Q)$ for ciphertext	20 – 1263	1263 bits
Max multiplicative depth L	N/A	31

Recall that FHE has a trade-off between implementation cost and supported complexity of computation: we can increase the multiplicative depth L and the plaintext modulus p^r by taking sufficiently high N and q , but this makes the homomorphic operations inherently slower. A typical range for the ring dimension N , still offering sufficient flexibility, is between 2^{14} and 2^{16} . BASALISC settles on a maximum value of $N = 2^{16}$. This allows us to choose ciphertext moduli up to $q = 2^{1782}$ at 128-bit security level. We get an acceptable number of multiplicative levels, even at a high-precision plaintext space (e.g., 31 levels at plaintext modulus $p^r = 127^3$ without bootstrapping; with bootstrapping, we get an arbitrary number of levels).

Table 1 shows the full parameter range supported by BASALISC and an example parameter set. Note that the largest ciphertext modulus is denoted by Q , but key switching matrices use an even larger modulus QP . Concretely, our largest supported modulus is $QP = 2^{1782}$ (limited by the 128-bit security target). The smallest supported modulus is $6 \cdot 2^{17} + 1$ – the smallest prime congruent to 1 modulo 2^{17} (see explanation in section 6.4).

3 DATA REPRESENTATION & ALGORITHMS

Homomorphic operations rely on arithmetic in the ring \mathcal{R}_q , which can be implemented efficiently based on the Chinese Remainder Theorem (CRT) [21]. For this purpose, we assume that the ciphertext modulus is given by $q = q_1 \cdot \dots \cdot q_k$, where the factors are distinct prime numbers satisfying $q_i = 1 \pmod{2N}$. Section 3.1 explains the operations in the coefficient ring \mathbb{Z}_q . Section 3.2 extends this to polynomials modulo $X^N + 1$.

3.1 Residue Number System

Arithmetic in \mathcal{R}_q can be split into many smaller rings \mathcal{R}_{q_i} simply by applying the Chinese remainder theorem. This idea is used commonly, and referred to as a Residue Number System (RNS). It brings an asymptotic speedup factor of $\mathcal{O}(k)$, but also simplifies the architecture since each q_i can be of size around 32 bits (compared to more than 1000 bits for q).

The restriction $q_i = 1 \pmod{2N}$ that we introduced above comes from the requirement that enables the NTT (see next section). This puts a lower bound of 17 bits on the size of each q_i since our design employs a maximum value of $N = 2^{16}$. Coupled with the requirement to have a sufficient amount of prime moduli to reach $\log_2(QP)$ of 1782 bits, we need q_i of at least 26 bits. However, we settle on 32-bit moduli, because it gives a better utilization for our on-chip memory buffer and simplified interaction with external

More information about key switching is given in Appendix A.

memory. Furthermore, we find that both 26-bit and 32-bit moduli result in the same complement of arithmetic units within our silicon area budget. For the example parameter set of Table 1, Q is a product of 42 primes and P is a product of 14 additional primes.

3.2 Number-Theoretic Transform

One of the most complex operations in FHE ciphertext arithmetic is polynomial multiplication. With a naive schoolbook method, multiplying two polynomials requires $\mathcal{O}(N^2)$ operations. For the large polynomial sizes innate to FHE, it is beneficial to resort to techniques based on the Fast Fourier Transform (FFT), allowing polynomial multiplications to be computed with $\mathcal{O}(N \log(N))$ operations only. The Number-Theoretic Transform (NTT) is the generalization of the FFT to finite fields. The NTT allows to use exact integer arithmetic, preventing round-off errors that are typical in real-valued FFT computations. The N -point NTT is given by

$$X[k] = \sum_{n=0}^{N-1} x[n] \omega_N^{nk} \pmod{q_i},$$

where ω_N denotes an N -th primitive root of unity. It can be shown that this root exists if and only if the modulus q_i is of the special shape $q_i = 1 \pmod{N}$.

In what follows, it will be useful to resort to the generalized description of the Cooley-Tukey algorithm that recursively re-expresses an NTT of size $N = N_1 N_2$ as N_2 inner NTTs of size N_1 and N_1 outer NTTs of size N_2 . Before the outer NTT, each output of the inner NTT is multiplied by a twiddle factor:

$$X[k_1 + N_1 k_2] = \sum_{n_2=0}^{N_2-1} \left(\sum_{n_1=0}^{N_1-1} x[N_2 n_1 + n_2] \omega_{N_1}^{n_1 k_1} \right) \omega_N^{n_2 k_1} \omega_{N_2}^{n_2 k_2}. \quad (2)$$

Radix- 2^k NTT algorithms can be obtained from this generalized description by choosing $N_1 = 2^k$ at each recursive decomposition. For example, by choosing $N_1 = 2$ and $N_2 = N/2$ or vice-versa, the well-known radix-2 Decimation-In-Time (DIT) and Decimation-In-Frequency (DIF) algorithms are obtained, respectively.

The NTT can be used for fast cyclic convolutions (polynomial multiplication modulo $X^N - 1$) through the convolution theorem:

$$ab \pmod{X^N - 1} = NTT^{-1}(NTT(a) \odot NTT(b)),$$

where \odot denotes point-wise multiplication. However, the rings used by BGV (and other FHE schemes) require negacyclic convolutions (polynomial multiplication modulo $X^N + 1$). If $q_i = 1 \pmod{2N}$, such that there exists a $2N$ -th primitive root of unity $\phi = \omega_{2N}$, then it can be shown that [2]:

$$\hat{a} \hat{b} \pmod{X^N + 1} = NTT^{-1}(NTT(\hat{a}) \odot NTT(\hat{b})), \quad (3)$$

where

$$\begin{aligned} \hat{a} &= (a_0, \phi a_1, \dots, \phi^{N-1} a_{N-1}), \\ \hat{b} &= (b_0, \phi b_1, \dots, \phi^{N-1} b_{N-1}). \end{aligned}$$

Thus, negacyclic convolutions can be computed using a regular NTT, together with a pre-multiplication and post-multiplication step with appropriate twiddle factors.

We note that the NTT can also be interpreted in terms of the Chinese Remainder Theorem, similarly to RNS. As a result, the

combination of using RNS for fast arithmetic modulo q and the NTT for fast polynomial arithmetic modulo $X^N + 1$, is often referred to as *Double-CRT*. Auxiliary homomorphic operations such as modulus switching and key switching rely on non-arithmetic operations that are not directly possible in Double-CRT format. Converting into and out of this format makes the auxiliary operations much more expensive than the basic ones: in practice, key switching dominates overall computation; it is roughly $100\times$ as expensive as multiplication.

3.3 Montgomery-Friendly Bootstrapping

As a common optimization, our Montgomery multipliers [32] are restricted to moduli of the shape $q_i = 1 \pmod{2N}$ that enable the NTT [31]. We refer to section 6.4 for a more extended explanation about the design of our multipliers. However, restricting the moduli in this way turns out to be incompatible with all currently existing bootstrapping methods for BGV [19, 23].

Consider for example the bootstrapping routine as implemented in the HElib library [23]. Let the plaintext modulus be $t = p^r$, then bootstrapping evaluates an adapted version of Equation 1 under the ciphertext modulus $q = p^e + 1$ that is significantly smaller than Q . It also involves an *exact* division by p^r , which is implemented based on arithmetic modulo p^r . However, both $p^e + 1$ and p^r are not Montgomery-friendly in general, so this cannot be done with our optimized Montgomery multipliers.

We propose a generalized version of bootstrapping that works with Montgomery-friendly primes exclusively. Our algorithm is simpler than all current approaches, yet it can be evaluated at exactly the same computational cost. The root of our algorithm is a new decryption formula that has sufficient degrees of freedom to take q as a product of Montgomery-friendly primes and does not involve an exact division operation. Consider the following lemma.

LEMMA 3.1. *Let $p > 1$ be a prime number, and let $e > r \geq 1$ and $q = 1 \pmod{p^e}$ be sufficiently high parameters. If (c_0, c_1) is a BGV encryption of m with plaintext modulus p^r and ciphertext modulus q , then we can decrypt it by computing*

$$c'_i \leftarrow [p^{e-r} c_i]_q, \quad w \leftarrow [c'_0 + c'_1 \cdot s]_{p^e} \quad \text{and} \quad m \leftarrow [\lfloor w/p^{e-r} \rfloor]_{p^r}.$$

Here we use $\lfloor \cdot \rfloor$ for coefficient-wise rounding to the nearest integer.

The first and second step in Lemma 3.1 are implemented based on the techniques of Bajard et al. [5]. The third step is identical to the bootstrapping algorithm from HElib [23]. More details such as the proof and pseudocode are deferred to Appendix B.

4 BASALISC ARCHITECTURE

The BASALISC architecture defines a semi-autonomous *co-processor* that accompanies and is managed by a commercial CPU. The CPU typically uses direct memory access (DMA) capability to transfer data to and from BASALISC’s memory subsystem, and uses either DMA or programmed IO to issue streams of FHE instructions to BASALISC. A simple interrupt-driven communication protocol allows BASALISC to indicate changes in state to the managing CPU, for example when a current instruction stream has been exhausted and BASALISC is awaiting further instructions.

These special moduli are called Montgomery-friendly primes in the literature [4].

BASALISC is an adapted Reduced Instruction Set Computer (RISC) architecture that allows for reasoning at diverse levels of executive abstraction. This multi-level approach aids in assuring the correctness of our system. Having a hierarchy of multiple intermediate representations and instruction sets, each with well-defined semantics, means that we can implement and test each stage of the compiler toolchain separately. In addition, different instruction set abstractions allow programmers to work at a higher level of abstraction while allowing compiler writers and library authors to reason about lower-level details such as scheduling and optimizations easily. For example, when writing a program to run on BASALISC, the programmer need not know about low-level data representations.

Specifically, we generate and reason over three distinct levels of instruction and typesystem abstraction.

- **Macro-instructions** are at the highest level, with the largest data types and the most complex operations. Entire ciphertexts, plaintexts, and key switching matrices are treated as basic data types at this abstraction level. Operations that can be described at this level include ciphertext addition, multiplication, automorphisms, modulus switching, key switching, and bootstrapping. Details about data representation and algorithms that implement those operations are opaque at this level of abstraction.
- **Mid-level instructions** expose the Double-CRT data representation used in BASALISC. The basic data type at this level is a residue polynomial (a polynomial in RNS representation) comprising up to 2^{16} 32-bit polynomial coefficients. Basic operations on these data types include pointwise modular addition and multiplication on vectors of coefficients; automorphisms; NTTs; and multiply-accumulate iterations commonly used in key switching. Also included in this list are memory management instructions that Load and Store data to and from the off-chip memory complement.
- **Micro-instructions** correspond very closely with the specific operations performed by the processing elements (PEs) of BASALISC. The basic data type at this level contains as many coefficient words (1024 or 2048) as can be processed simultaneously by a PE or accessed in one on-chip memory cycle. Instructions at this level are delivered via the Peripheral Component Interconnect Express (PCIe) interface to the BASALISC processor for execution. This instruction level also includes rudimentary machine control instructions.

Table 2 shows examples of operation code mnemonics (opcodes) from each of our instruction sets.

FHE programs are deterministic: data dependencies do not exist since variables are encrypted, and hence all branches must be translated to some form of predicated execution. BASALISC takes advantage of this determinism: memory allocation is bound at compile time, and the size of all data operands is also bound at compile time. As a result, BASALISC uses a register-like addressing mode for all levels in its memory hierarchy, and has no need for cache-like structures that bind allocation of memory resources at runtime.

As shown in Table 2, BASALISC uses explicit Load-Store-Move semantics for managing its memory hierarchy, which is divided into

Table 2: BASALISC Example Opcodes. We omit operand specifiers.

ISA	Opcode	Semantics
Macro	LOAD	Move data from distant to near memory
	KSW	Key switch a ciphertext
	MORPH	Perform automorphism on a ciphertext
Mid	MULI	Element-wise multiply a residue polynomial by constant
	NTT	Compute NTT of residue polynomial
	FBE	Fast Base Extension
	HCF	Halt and await new instructions
Micro	NTT1	Perform an iteration of the first pass of an NTT
	MAC	Multiply two operands, add result to accumulator

Table 3: BASALISC Micro-level operand addressing modes.

Mode	Definition
\$XXX	address of chunk in distant memory, used only for LOAD/STORE
rXXX	address of chunk in middle memory
tXX	register number in near memory
nXXX	immediate 32-bit scalar
iXXX	index into table of prime moduli

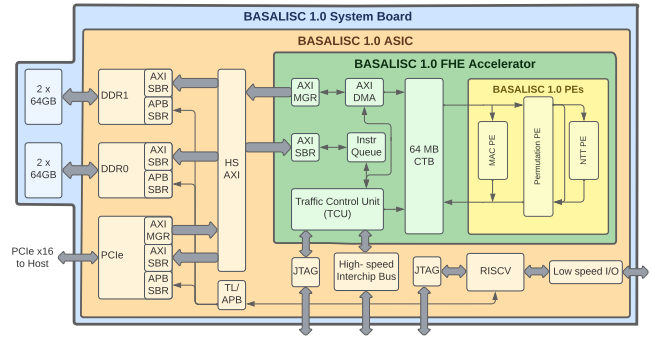
four levels. *Distant* memory is accessible only by LOAD and STORE explicit references, as well as by DMA from host memory. In our current single-chip design, distant memory is realized in off-chip Dynamic Random Access Memory (DRAM) to allow for sufficient storage for a meaningful working set of data. *Middle* memory is accessible by LOAD and STORE operations, and also by direct addressing in data processing instructions. Middle memory is realized in on-chip Static Random Access Memory (SRAM) in our current design. *Near* memory consists of a register set integrated in our multiply-accumulate functional unit, allowing for efficient addition, multiplication, and multiply-accumulate operations common for example in key switching inner loops, without reaching back to Middle memory. Finally, immediate scalar 32-bit values may be specified directly in the instruction stream for some instructions such as *Multiply by Immediate* (MULI) (Table 2).

Operand specifiers exist at each level of the BASALISC ISA. Table 3 shows examples of the addressing modes for operand specifiers at the Micro-instruction level. At this level, operands are either “chunks” of 2048 32-bit coefficients within a residue polynomial, 32-bit scalar values, or natural number indices into tables of moduli.

5 BASALISC HARDWARE DESIGN

Figure 2 shows a block diagram of BASALISC 1.0 – the first implementation in the BASALISC architecture family. BASALISC 1.0 is a single-chip FHE coprocessor, designed in a 12nm Global Foundries process, with additional off-chip memory; high-speed connectivity to its host system; and extensibility via a high-speed inter-chip interconnect. In contrast to other FHE hardware accelerators, BASALISC 1.0 reduces cost and manufacturing risk by relying only on commercially available standard packaging, DRAM, and PCIe technologies.

BASALISC System Board. At left in Figure 2, we instantiate *distant* memory using two Double Data Rate 4 Synchronous Dynamic Random-Access Memory (DDR4)-3200 subsystems, each providing up to 128 GB of DRAM and 25.6 GB/s of bandwidth. At bottom left of the diagram is the 26 GB/s (near-peak) PCIe x16 channel that connects BASALISC to its host and carries data and instructions.

**Figure 2: BASALISC 1.0 System Diagram.**

For many applications, our on-chip 64 MB *middle* memory SRAM array that we call the Ciphertext Buffer (CTB) is too small to hold the sizeable working sets of ciphertexts and key switching matrices. Thus, the CTB will suffer high capacity miss rates. Therefore, we expect these applications to be performance-limited by our twin DDR4 channel bandwidth. When the 256 GB complement of DDR4 is too small, performance will be limited by PCIe’s long latency and low bandwidth to host memory.

BASALISC ASIC. The orange rectangle at center in the diagram denotes the logical boundaries of the BASALISC 1.0 ASIC. Shown at left within that box are the controllers and physical interfaces (PHYs) for DDR4 and PCIe. These PHYs connect to the 512-bit wide Advanced eXtensible Interface 4 (AXI4) interconnect that transfers data between the DDR, PCIe, and the CTB. Both the AXI4 and CTB operate at a target cycle time of 1 GHz. As a result, the AXI has a peak bandwidth of 32 GB/s for each endpoint connection, all running in parallel. At bottom right in the diagram is the RISC-V CPU core we use to configure BASALISC at startup time. During normal operation, this CPU core is inactive, so we do not discuss it further. Also shown are the Joint Test Action Group (JTAG) external I/O connections used in testing and debugging BASALISC 1.0.

BASALISC FHE Core Processor. In green is the core BASALISC 1.0 FHE Accelerator. This subsystem includes the CTB, AXI4 infrastructure, the instruction queue, and a Traffic Control Unit (TCU) that manages instruction execution in the system. Because control flow such as branching and iteration is not needed in BASALISC, the TCU is much simpler than in a traditional CPU.

The CTB is a single-port SRAM array that can either read or write 2048 32-bit residue polynomial coefficients every machine cycle, providing a total bandwidth of 8 Tbps (at 1 GHz operation) to our complement of data Processing Elements (PEs) shown in yellow.

As an advantage of the compile-time determinism of the FHE programming model, the BASALISC CTB comprises an addressable set of ciphertext registers, instead of requiring the functionality of a cache memory. This set of registers is compiler-managed with a true least-recently used (LRU) replacement policy. CTB bandwidth is not materially affected by concurrent transfer between distant memory and the CTB: roughly at most 0.3% of CTB access cycles are used by our total distant memory bandwidth.

6 BASALISC PROCESSING ELEMENTS

BASALISC 1.0’s on-chip PEs and their connection to the CTB are shown on the far right in Figure 2. The BASALISC PEs that rely on the CTB for data are the Multiply-Accumulate (MAC) PE (used in ciphertext addition, multiplication, and for kernels of operations such as key switching); the Permutation PE (used to permute data into preferred orders to achieve NTT processing, and also used for automorphisms); and the NTT PE (used to accomplish number-theoretic transforms efficiently). We describe each of their capabilities below.

Whereas Figure 2 shows single PE instances, their implementation is a massively multicore architecture that exploits innate parallelism in FHE ciphertext computations. FHE arithmetic in RNS representation offers four types of parallelism: (i) over multiple ciphertexts, (ii) over the polynomials within a ciphertext, (iii) over the residue levels of a polynomial, and (iv) over the coefficients of a residue polynomial. Prior work has focused on (iii), instantiating multiple so-called Residue Polynomial Arithmetic Units (RPAUs) [30, 40, 49]. In contrast, BASALISC focuses on exploiting (iv), due to two key observations. First, the number of residues decreases with the modulus level in the BGV scheme, leading to would-be idle RPAUs as the computation gets closer to bootstrapping. Second, as the lowest level of parallelism, coefficient-level parallelism offers the best opportunity to exploit locality of reference and limit CTB thrashing.

6.1 Number-Theoretic Transform PE

Because of the focus on coefficient-level parallelism, BASALISC implements a high-radix NTT PE. We expect that many BASALISC FHE applications will employ ring dimension $N = 65536 = 256^2$ to enable bootstrapping and thus arbitrary-depth computation. Thus, our NTT PE employs a radix-256 butterfly, allowing us to compute 65536-point NTTs with only two round trips to memory for each coefficient. NTTs of smaller sizes can be computed through shortcut paths in our NTT butterfly network.

Following the generalized Cooley-Tukey NTT description of Equation 2, a radix-256 NTT chooses $N_1 = N_2 = 256$. The main arithmetic NTT unit consists of a 256-point NTT (that computes the inner N_1 -point NTT and outer N_2 -point NTT) followed by 255 post-multipliers (that multiply with the twiddles $\omega_N^{n_2 k_1}$). We employ a standard DIF flow graph for the 256-point NTT, where we replace multiplications $\omega^0 = 1$ with simple pipeline balancing registers. Through this optimization, the inner NTT is implemented with only 769 modular multipliers, instead of $N/2 \log(N) = 1024$.

As shown in Equation 3, an additional pre-multiplication and post-multiplication step is required to construct negacyclic forward and inverse NTTs from regular NTTs. Because a radix-256 butterfly already includes an array of 255 post-multipliers, it suffices to add 255 pre-multipliers to efficiently support negacyclic NTTs. The result is a 3-stage NTT architecture, as illustrated in Figure 3 for a scaled-down radix-4 unit. In Figure 4, we illustrate how a radix-4 unit is composed to compute the full NTT flow graph in two passes that each take 4 chunks. In between the passes is an implicit memory transposition that we enable with a *conflict-free CTB design*.

Our NTT PE instantiates four parallel 3-stage NTT units. Each unit is deeply pipelined with 40 pipeline stages in order to run at 2 GHz. Together, these four parallel pipes consume 1024 32-bit




Figure 3: Radix-4 negacyclic NTT unit with pre- and post-multiplier arrays.




Figure 4: 16-point radix-4 negacyclic NTT flow graph. Extra negacyclic twiddles (in blue) are decomposed into two pre-multiply passes.

residue polynomial coefficients at that 2 GHz rate – sufficient to consume all available data bandwidth from the CTB.

6.1.1 Conflict-Free Schedule. A well-known performance inhibitor for NTTs is that successive NTT passes access coefficients at different memory strides, introducing access conflicts in memory. Prior NTT accelerators present custom access patterns and re-ordering techniques that only work for small-radix NTT architectures [37, 41] or require expensive in-memory transpositions [42]. BASALISC avoids reinventing the wheel, instead building upon years of DSP literature [14, 24, 29]. The most high-performance FFT accelerators present *conflict-free schedules* [36, 38, 48] to tackle this exact issue.

Conceptually, a $N = N_1 N_2 = 256^2$ -point radix-256 NTT can be represented as a two-dimensional NTT, where the data is laid-out with $N_1 = 256$ rows and $N_2 = 256$ columns. In this format, the inner N_1 -point NTT requires coefficients in column-major order, whereas the outer N_2 -point NTT requires data in row-major order. The crux of building conflict-free NTT schedules is to structure the data so that it can be read out in either order without bank conflicts. This requires a minimum of 256 independently addressable banks, each containing 2^{16} bank addresses (for a total CTB size of 2^{24} values).

We employ a conflict-free layout based on XOR-permutations [38], as illustrated in Figure 5. In this layout, data with logical address $\{row, col\}$ is stored at $bank = row \oplus col$. This layout ensures that each unique index in every row and column corresponds to a unique physically accessible bank of CTB SRAM.




Figure 5: Example conflict-free CTB layout for a 16-point radix-4 NTT. Data is striped using the equation $bank = row \oplus col$, which ensures that both entire columns or entire rows can be read out without bank conflicts. The on-the-fly Permutation PE maps values from bank order into natural order, as illustrated for access to the second column.

When reading rows or columns from the CTB, values come out of memory in *bank order*, one value for each bank from bank 0 to 255. However, operations like NTT require values in *natural order*: when accessing a row, we need values sorted by column from 0 to 255, and when accessing a column, we need values sorted by row from 0 to 255. Thus, when accessing row r , we must map bank i to index $i \oplus r$. Likewise, when accessing column c , we must map bank i to index $i \oplus c$.

We build a custom “on-the-fly” Permutation PE to compute these XOR-based permutations as data moves to or from the other PEs. Furthermore, we are the first to observe a remarkable optimization opportunity for this unit. By implementing a slightly more general permutation PE that supports permutations of the form $i \mapsto (i \cdot a + b) \oplus c$, we can not only use the Permutation PE to implement conflict-free XOR permutations, but also any BGV ring automorphism *without additional hardware*. The Permutation PE is described in more detail in section 6.2.

6.1.2 Twiddle Factor Factory. Similarly to polynomial residue coefficients, twiddle factors in BASALISC are 32-bit integers. There are N twiddle factors for each residue for both forward and inverse NTT, and a maximum of 56 residues at max-capacity key switching, together requiring ~ 29.4 MB of twiddle factor material in a naive implementation. Moreover, our four NTT units have 5116 multipliers total that must be fed each cycle with twiddles, requiring massively parallel access into this storage memory. BASALISC prevents this storage requirement in two ways. First, we contribute new insights and a twiddle decomposition method, that reduces the required parallel number of distinct twiddle accesses. Second, we develop a custom *twiddle factor factory* that drastically reduces the number of twiddles stored. In the remainder, we analyze only the forward NTT, but note that identical optimizations apply to the inverted twiddles for the inverse NTT.

For a forward negacyclic NTT, each input x_i is pre-multiplied by the twiddle $\phi^i = \omega_{2N}^i$. Using techniques from the DSP literature [17], we propose to decompose the additional negacyclic twiddles to extract a regular pattern, and to distribute them evenly between the two NTT passes in the flow graph. This is illustrated in Figure 4 by the extra twiddles present in blue. The benefit of this technique is twofold. Firstly, it can be easily seen that through this technique, the pre-multiplications become identical for each chunk in both passes.

This allows the four NTT units to share the same pre-multiply twiddles, and drastically reduces the total number of pre-multiply twiddles from $N = 256^2$ to $2 \cdot \sqrt{N}$, easily fitting in a small SRAM. Second, the internal butterfly twiddles (powers of ω_{256}) are now a strict subset of the pre-multiply twiddle in the first pass (powers of ω_{512}). Both can therefore be routed from the same small SRAM.

The remaining twiddle factor complexity sits in the post-multiply twiddles. For each chunk k , there are 255 twiddles $\omega_{256^2}^{ik}$. An SRAM storing vectors of 255 twiddles with depth 255 for each residue is still much too large. We propose a technique to reduce the *width* of this SRAM. It can be coupled with techniques that reduce the *depth* of this SRAM, such as On-the-fly-Twiddling (OT) [26]. To reduce the width, we propose a *power generator circuit* that trades SRAM storage for multipliers. The main idea is as follows. By using the identity $\omega_{256^2}^{ik} = \omega_{256^2/k}^i$, it can be observed that the required twiddles for chunk k are always the 255 consecutive powers of a *seed* value $\omega = \omega_{256^2/k}$. Using only ω , we can compute its successive powers in a number of multiply layers. The first layer computes ω^2 from ω , with a single multiplier. The second layer takes ω^2 and ω to compute ω^4 and ω^3 , and so forth. Every multiplier in the circuit produces a unique value that is used as an output, so the number of multipliers to generate 255 powers from ω is simply 254. Using this technique, instead of storing vectors of 255 twiddles with depth 255 for each residue, it suffices to store the single seeds with depth 255.

6.1.3 Related Work. Many NTT architectures have been proposed in the literature. The architecture that is closest to ours is that of F1 [42], another concurrently-developed BGV accelerator. Targeting a smaller parameter set than BASALISC, F1 similarly implements a radix-128 butterfly to efficiently support $N = 16384 = 128^2$ -point NTTs.

To allow both row-major order and column-major order NTT passes, F1 implements an explicit matrix transpose unit. In contrast to BASALISC’s conflict-free layout, this unit requires an in-memory transposition that will stall the NTT pipeline. To prevent these stalls, F1 fully pipelines their transpose PE. This requires the transpose PE to implement an SRAM that buffers nearly the full polynomial output of an NTT pass, which is much more expensive than BASALISC’s cheap Omega-Network.

F1 adopts prior techniques to merge the negacyclic pre-multipliers into a DIT NTT flow graph [34, 41], an “optimization” that is represented in many works as going from $N + N/2 \log_2(N)$ to only $N/2 \log_2(N)$ multiplications. The salient assumption for these formulas is that an $N = 256$ -point NTT datapath has a DIT butterfly at every node, for a total of $N/2 \log_2(N) = 1024$ multipliers. On the other hand, by eliminating those that multiply with $\omega^0 = 1$, a 256-point NTT requires only 769 multipliers. Together with the pre-multipliers, both NTT implementations have an identical number of multipliers. This technique is therefore not an optimization to reduce multiplier counts, but rather a way to distribute multipliers more homogeneously in the flow graph. We do not adopt it in BASALISC, because it removes the fixed static twiddle pattern within the 256-point NTT that we exploit heavily in our twiddle factor factory unit. F1 does not describe how they implement or simplify their large twiddle factor SRAM.




Figure 6: Simplified Diagram of MAC PE architecture.

6.2 Permutation PE

A pair of Permutation PEs forms the interface between the CTB and the other PEs. We observe for the first time that a slightly more generalized Permutation PE can support both conflict-free schedules required by NTT operations, as well as BGV automorphisms *with the same hardware*. In order to do so, the Permutation PE is generalized to compute permutations of the form $i \mapsto (i \cdot a + b) \oplus c$. Each permutation unit reorders an array of input coefficients to produce a permuted output array of the same length.

The *Read Permutation PE* unscrambles data in conflict-free CTB bank ordering in order to pass it to the other PEs expecting natural ordering. It is a specialized instance of the more general Permutation PE that only implements permutations $i \mapsto i \oplus c$, requiring values $a = 1$ and $b = 0$. The *Write Permutation PE* passes data in the opposite direction. It implements the general permutation $i \mapsto (i \cdot a + b) \oplus c$ in order to re-scramble the data into its conflict-free layout, or to compute ring automorphisms. In the latter case, the output of the Read Permutation PE is fed directly into the input of the Write Permutation PE to achieve the complete operation of the automorphism.

Each Permutation PE itself is split into two portions. Firstly, the data-permutation portion of the logic is implemented using 2×2 switch nodes placed using an Omega-Network topology. Secondly, a configuration portion takes constants a , b , and c in order to generate the routing pattern for the switches in the network. The configuration portion of the logic attaches the routing pattern to the data and the combined payload word is sent through the network. The switch nodes forward the data according to the least significant bit of the pattern part of the payload data, which is also removed before forwarding. Thus the message is reduced by one bit at each stage of the network and at the end the payload only contains the data portion.

6.3 Multiply-Accumulate PE

We realize modular addition and multiplication for FHE in the Multiply-Accumulate PE (MAC), shown in Figure 6. This pipelined unit can start 2048 32-bit modular addition or modular multiply operations each cycle, if data is available. Because the MAC PE is built with asynchronous logic, it free-runs at 1.6 GHz when not accessing the 1 GHz CTB. Therefore, operations that read and/or write to the CTB are limited by the 1 GHz CTB bottleneck, while other operations that operate on local data (accumulator register or register file) can accelerate to 1.6 GHz, without using any additional logic. The asynchronous logic provides significant area and latency savings over implementing wide Clock Domain Crossings (CDCs)

Table 4: Area and power comparison of NTT single-butterfly unit with original and optimized Montgomery multipliers at 1 GHz.

Multiplier Design	Area	TDP @0.72V, 125C
Unoptimized	3768 μm^2	7.2 μW
Optimized	2052 μm^2	4.3 μW

instead to achieve this 60% performance using a clocked approach. At left in the figure, the 2048 a inputs, each 32-bit in size, come from the CTB. The b inputs are replicated copies of a 32-bit constant from the instruction stream. The MAC includes a 16-entry Register File (RF), shown at top in the figure. In addition, there is a single accumulator register at the output of the adder/subtractor/accumulator unit, shown at right in the figure.

Using the multiplexers shown in the figure, this arrangement can accomplish a variety of functions. Residue chunk multiplication by or addition of a constant to each coefficient can be accomplished at full rate: 2048 32-bit operations per cycle. Multiplication or addition of chunks when both are sourced directly from the CTB can be accomplished at half-rate, using a register to buffer one operand from the CTB, and directly feeding the second operand into the operation from the CTB in a second read cycle. Acceleration of tight kernels that repeatedly process the same chunks can be achieved by storing up to 16 different chunks in the RF and then operating on them at full rate. Finally, the MAC has a multiply-add capability similar to that often found in digital signal processors, allowing double-rate processing: a multiply and accumulate in every cycle. The above possibilities are impacted by the write bandwidth needed to the CTB for results. Write operations might occur as often as for every chunk result, or much less often when the local RF or the accumulator are used to store results during tight kernel operations.

A particularly important example of kernel acceleration in the MAC is key switching from Appendix A. We expect key switching to make up the large majority of the workload of a typical FHE program. The inner loop of our key switching algorithm is a “fast base extension” subroutine that pre-computes a table of about 12 residue polynomials, and then computes many (around 40) different weighted sums of those twelve values, with constant weights. Use of the local registers in the MAC PE and the compound multiply-accumulate function realizes a $44\times$ improvement compared to a naive design. In addition, this approach reduces use of the CTB during fast base extensions to 10.6% versus nearly 100%, saving 90% of the CTB for use by the other PEs.

6.4 Modular Multiplier Arithmetic Optimization

Both the MAC and NTT Butterfly units use Montgomery modular arithmetic, optimized for Montgomery-friendly primes [31], and matched to our novel bootstrapping approach. Specifically, instead of supporting the full 32-bit prime value, the multiplier is optimized to only support a subset compatible with our approach, where the lower 17 bits of the prime are fixed (bits 16:1 are tied to 0 and bit 0 is tied to 1). This optimization of the Montgomery modular arithmetic saves 46% in area compared to a generic Montgomery multiplier that can support all moduli. The results are summarized in Table 4.

6.5 Memory Subsystem

BASALISC includes a four-layer memory hierarchy for storing ciphertexts and keys. From farthest to nearest to the PEs, these are the distant memory (off-chip DRAM), middle memory (CTB), the MAC RF, and the accumulator register, as shown in Figure 2. The capacity and latency of each level in the hierarchy are shown in Table 5.

As shown in the table, the layers in our memory hierarchy exhibit diverse latencies and capacities typical of computer memory hierarchies, where lower latency layers have smaller capacities. A significant difference between typical memory hierarchies and that of BASALISC is the working set size that each layer can hold. Nevertheless, we expect capacity limits of layers in our memory hierarchy to be a major limiter of system performance. In particular, we expect minimal locality of reference for key switching matrices, each of which is larger than the entire CTB.

Table 5: Memory hierarchy for ciphertext and key storage.

Memory	Capacity	Round-trip latency
Off-chip DDR	256 GB	>100 ns
CTB	64 MB	~3 ns
MAC RF	128 kB	~1.25 ns
MAC ACC	8 kB	0.625 ns

6.5.1 Middle Memory - the CTB. The 64 MB CTB contains 2^{24} locations, each of which holds a 32-bit residue polynomial coefficient. In our largest supported parameter set, a single residue polynomial consists of $N = 2^{16} = 64K$ coefficients and occupies one entire page of the CTB. As explained before, a residue polynomial is arranged conceptually as a 256-by-256 rectangular array in the CTB, and coefficients are physically striped to support the conflict-free NTT schedule. For smaller ring dimensions, polynomials are arranged as a 256-by- $(N/256)$ array, and a single CTB page will contain multiple residue polynomials.

6.5.2 Distant Memory - the DRAM array. The DRAM system, as described in the interface section above, comprises two independent DDR4 interfaces in parallel, each supporting up to 128 GB of DRAM. The DRAM serves as the staging area for data that is scheduled for processing and for results that are ready for retrieval by the host computer. The two interfaces allow us to maximize the practical throughput of the DRAM subsystem, by avoiding collisions between the PCIe-to-DRAM and FHE-to-DRAM access streams. The PCIe can access data on one DRAM interface without interfering with the FHE accelerator. At the same time, the FHE accelerator can process data using the other DRAM interface. Since the dataflow will be known in advance by the compiler, the data can be arranged in such a way between the two DRAM instances as to minimize the interactions between the two.

6.5.3 The Instruction Buffer. The BASALISC instruction buffer is organized as a batch queue, and is maintained by the TCU. Each instruction takes tens of cycles to execute and requires a substantial amount of memory in the CTB, therefore the instruction queue will be fairly short (128 to 1024 instructions, depending on our performance analysis). The queue is loaded periodically with new




Figure 7: BASALISC Software Toolchain.

instructions when the host knows that there is a small amount of pending instructions in memory left. Since the compiler knows in advance about the flow of instructions and data in memory, this can be predicted and is the reason why there is little need for a long program queue.

7 BASALISC COMPILATION AND SIMULATION TOOLS

Figure 7 shows the main components, languages, and intermediate data representations in the BASALISC software toolchain. The dashed boxes in the figure represent our two main software tools: Artemidorus is our compiler, which takes input programs written in a Domain-Specific Language (DSL) and outputs one of our three distinct instruction sets. Simba is our simulator, which takes instruction traces as input, and produces either a performance report or concrete result values as output.

Artemidorus. As shown at top left in the figure, our toolchain begins with high-level DSL that allows programmers to create FHE applications for BASALISC to execute, and which features data types including fixed-point numbers, vectors, and matrices. The program passes through several stages in Artemidorus. Bootstrap operations, vector operations, and matrix operations are expanded into BGV primitives; key switching operations are inferred; each operation is tagged with the length of its modulus chain, and then expanded into primitive operations on individual residue polynomials for each factor in the modulus; and finally memory regions and registers are allocated and instructions are scheduled.

Artemidorus produces instruction traces at our three levels of the ISA, which pass to Simba for performance or correctness simulation. Especially *Simba-micro*, the micro-level performance simulator shown on the bottom right, presents an integral part of evaluating BASALISC at this point in the design stage. We now describe it in detail.

Simba-micro. Our micro-level performance simulator employs a step-based operational semantics to model the execution of the BASALISC coprocessor. There are five basic operational components: the CTB, and the four PEs (MAC, Read Permutation, NTT,

and Write Permutation). Each of these components operates at a different internal frequency (Table 6). The simulator models a micro-instruction’s life cycle from instruction dispatch, to data transfer from CTB to the appropriate functional unit, to proceeding down the pipeline, to the “writeback” phase.

In order to account for the different clock rates of the different components, we use a global “micro-clock” which operates at 6 GHz as the time increment for the model’s step function. We made the simplifying assumption that the MAC operates at 1.5 GHz. In this way, we were able to model each PE’s progress by causing the CTB to be accessible every 6 micro-cycles, the MAC every 4 micro-cycles, and the permutation/NTT units every 3 micro-cycles. This behavior is modeled by supplying each component with a wait counter which is reset to these values every time it is accessed; the component is only accessible if the counter is 0. If a component is accessible but has no work to do in the given micro-cycle, it simply waits until it has something to do.

Each individual PE is modeled as a pipeline with a certain number of stages and “stage capacity” (number of coefficients that fit in each pipeline stage). The MAC’s stage capacity is 2048 coefficients, while the other three have a capacity of 1024. Every time the given PE is enabled (its wait counter is 0), the pipeline advances. When there is a write at the end of the pipeline, it stays at the end of the pipeline until the CTB is available for writing.

Oftentimes, the CTB can be used for either reading data or writing data in a given micro-cycle. This occurs whenever the next instruction reads from the CTB, while there is a write “waiting” at the end of either the MAC or write permutation pipeline (or both). In this scenario, we opted to always favor reads over writes; therefore in our simulations, pipelines tend to fill up. Once a pipeline is full, instruction dispatch is no longer possible to that pipeline, and the control mechanism allows the pending writes to occur. After execution, the following data is reported by Simba-micro: number of CTB cycles (i.e., 6 micro-cycles) of execution, overall CTB utilization (percentage of time spent reading/writing/stalling), and utilization of each PE (how “full” the pipelines are, broken down by % of time).

8 EVALUATION OF BASALISC

At present, BASALISC is an architecture with an implementation-in-progress but not delivered to silicon yet. We evaluate the architecture and design of BASALISC in diverse ways at this point in the design cycle.

8.1 Physical Realizability

One way to evaluate the design of BASALISC 1.0 and the BASALISC architecture that it implements is by undertaking a physical design implementation in a practical semiconductor process, with a reasonable target die size and operational frequency target.

One resulting evaluation criterion that can be objectively measured using this approach is timing closure – the verification that, with placement and routing of key blocks complete, and using industry best practice estimation of inter-block wire delays based on a mature floorplan, the design achieves a target operating frequency that yields useful levels of performance. In the case of BASALISC 1.0, we completed placement and routing of the novel circuitry - our

Table 6: Performance characteristics of BASALISC hardware elements.

Component	f_{max}	Area	TDP @0.72V	Throughput
CTB	1.0 GHz	77.9 mm ²	9 W	2 × 32 Tb/s
MAC PE [†]	1.6 GHz	7.17 mm ²	18.6 W	102 Tb/s
NTT PE	2.0 GHz	16 mm ²	24.6 W	32 Tb/s
Permutation PE	2.0 GHz	0.16 mm ²	~0 W	32 Tb/s
PCIe	500 MHz	12 mm ²	5 W	26 GB/s
DDR	800 MHz	18.2 mm ²		51 GB/s
Overall	>1.0 GHz	150 mm ²	57.5 - 115 W	N/A

[†] Operation of PEs above the frequency of the CTB is advantageous when they can run independently of the CTB.




Figure 8: Floorplan of BASALISC 1.0 with all cells placed and intra-block routing complete. Note that the MAC PE and Permutation PE are interleaved within the CTB.

PEs - and the CTB RAM block. Our operational frequency target was a minimum of 1.0 GHz at the standard “slow-slow” (SS) process corner and a supply voltage of 0.72V in the 12nm low-power Global Foundries process. We achieved timing closure for the diverse functional units at the frequencies given in Table 6. Our target die size is constrained to 150mm², with an aspect ratio of 2 : 1 or less. Our floorplan shown in Figure 8 uses actual IP block sizes for DDR4 DRAM, PCIe, our RAM array, and placed and routed PEs, providing a 14.4mm × 10.4mm die size that satisfies both of those constraints.

8.2 Logic Emulation & Formal Verification

A commonplace verification step prior to ASIC manufacturing provides yet another evaluation criterion: successful *hardware emulation* of critical logic in the design. In the case of BASALISC 1.0, that critical logic is the set of processing elements (MAC, permutation, and NNT PEs). We successfully emulated each PE in full, using test vectors extracted from our Verilog testbenches and our formal models of each PE. Each PE passed its emulation test vector suite.

In addition to hardware emulation, BASALISC employs formal methods with two primary goals: first, that the design be proven mathematically correct, and, second, that the design be proven consistent at every intermediate representation by demonstrating proof of equivalence. For both the mathematical and consistency proofs, BASALISC employs the Cryptol language [27] and related tools.

In order to satisfy mathematical correctness, top-level FHE algorithms are expressed as a mathematical model in Cryptol. Subsequently, using Cryptol’s proof capabilities, the mathematical model is proven to sustain a set of separately-developed correctness definitions. Proof of equivalence is provided through a two-step approach. First, formal equivalence is proven between the high-level mathematical Cryptol model and a low-level logic-oriented Cryptol description using SAW [8]. Next, the low-level Cryptol is converted to Verilog that we prove equivalent to the optimized implementation-Verilog using the commercial Synopsys Formality tool.

8.3 Benchmark Performance Simulation

We evaluate BASALISC 1.0 performance on different benchmarks: a micro-benchmark that measures a single bootstrapping operation, and a macro-benchmark comprising a single iteration of logistic regression training over encrypted data. In both benchmarks, we use the example parameter set from Table 1. Benchmarks for basic and auxiliary homomorphic operations are provided in Appendix C.

8.3.1 Micro-Benchmark. We estimate execution time for a single bootstrapping operation using our cycle-accurate simulator Simba-micro. We set the bootstrapping parameter from section 3.3 to $e = 4$. Simulation results show that bootstrapping consumes only 40ms of execution time. In comparison, HELib takes 160s to bootstrap a single ciphertext with comparable parameters on a Intel Xeon E5-2630 v2 CPU at 2.6 GHz running a single thread. Hence, we achieve a speedup of 4,000 times.

8.3.2 Macro-Benchmark. We estimate execution time on one iteration of secure logistic regression training. We apply the algorithm from Chen et al. [10] to a 1,024-sample, 10-feature infant mortality data set from the US Centers for Disease Control.

We apply three changes to the original algorithm: we replace the FV scheme by the BGV scheme; we replace sign extraction by an improved variant that has higher precision; and we replace the sigmoid function by the piece-wise linear approximation over $[-63, 64]$ that is shown in Figure 9. These three change were proposed in the DARPA DPRIVE program. Since the adapted algorithm is not important for our purpose, we do not explain it in detail here.

We use the example parameters from Table 1, resulting in ciphertext size 21 MB and key switching size 84 MB. The single

iteration of logistic regression training includes 513 bootstrapping operations. We note that this application actually uses a variant of bootstrapping that performs data scaling at the same time. This variant is much heavier than our micro-benchmark. As a BASALISC instruction trace, the logistic regression is composed of 908,660 high-level, 850,564,991 mid-level, and 27,251,778,560 micro-level instructions. Table 7 shows how the sigmoid is broken down into mid-level BASALISC instructions. Each sigmoid accounts for 2 of the 513 bootstrappings.




Figure 9: Sigmoid and PL approximation.

Table 7: Sigmoid mid-level instructions.

ADD	218,650
SUB	150
MUL	120,948
MORPH	3192
NTT	114,389
INNT	17,954
CONVERT	56,204
FBE	1,787

Again, we simulate the resulting micro-level trace using our cycle-accurate simulator Simba-micro. The trace consumes 40.5s of simulated execution time: 3,491 times slower than running the same algorithm on a single core Intel Xeon Silver 4210R CPU at 2.4 GHz *without* using FHE. Since HELib requires 160s to bootstrap a single ciphertext, the total running would be 82,080s accounting for the 513 bootstrapping occurrences.

Hence we achieve a speedup of *at least* 2,025 times. Note, however, that this is an underestimation, because logistic regression uses the heavier version of bootstrapping that is not present in HELib.

8.4 Related Work

As part of BASALISC’s evaluation, we attempt a comparison to prior FHE accelerators. This comparison is complicated, because many prior architectures do not report bootstrapping benchmarks or simply do not support it [9, 16, 30, 35, 37, 44–46, 50]. These architectures support only unrealistically small parameter sets, often allowing them to fully compute on-chip. Furthermore, not all accelerators implement full FHE computations, but rather individual computations such as the NTT. As one outcome, these other approaches require frequent interaction with a host processor to sequence operations and combine results. In this category, HEAWS [50] reports a 5.5× speedup compared to a software reference for a low-complexity neural network with multiplicative depth 4. HEAX [37] achieves more significant acceleration numbers, in the order of 200× for high-level CKKS operations such as key switching, compared to SEAL [43]. Another recent accelerator for CKKS, Medha [30], shows a 130× speedup for ciphertext multiplication. Compared to HEAX, Medha chooses to sacrifice throughput in order to optimize for latency. However, at larger parameter sets, FHE kernels and applications provide a multitude of parallelisms and possibilities for re-ordering to eliminate data dependencies. For BASALISC

applications, we expect very limited read/write dependency stalls where low-latency would be favorable.

The architecture that is closest to ours is F1 [42]. F1 is an ASIC architecture targeting the same die size (150mm²), technology node (12nm GF), and clock frequency (1 GHz), and it also implements bootstrapping. Whereas our micro-benchmark shows that BGV bootstrapping takes 40ms in BASALISC 1.0, F1 [42] reports a bootstrapping time of only 2.4ms. However, these numbers are not comparable: F1 provides lower security (their ring dimension N is 4 times lower) and supports a plaintext space of only 1 bit with no packing. BASALISC supports plaintext modulus 127³ with packing capability.

Finally, a recent study by De Castro et al. highlighted the memory bottleneck of FHE acceleration [15]. The starting point for their analysis is a CPU-like architecture, where ciphertexts do not fit in the Last-Level Cache (LLC). However, BASALISC’s compiler-managed on-chip CTB is very different from a typical LLC. Moreover, at 64 MB, we are able to fit several ciphertexts within the CTB. Nevertheless, we also observe in BASALISC that data movement of *key switching matrices* will often be the practical bottleneck of FHE applications.

9 CONCLUSION

FHE enables new privacy-preserving applications, but its adoption is limited because of high computational costs. BASALISC accelerates FHE computations by more than three orders of magnitude over CPU performance, and thereby takes a step toward practical feasibility.

In contrast to many prior works, BASALISC supports all BGV operations, including bootstrapping, in a single ASIC architecture. Our design includes a complete memory hierarchy, and an ISA that supports different levels of abstraction. Moreover, we propose several new hardware improvements: we implement a 32 Tb/s NTT architecture, and show that its permutation unit can be generalized to compute BGV automorphisms without additional area. We also save over 40% in area and power consumption by restricting our multipliers to Montgomery-friendly primes. We show that this optimization still allows BGV bootstrapping, and therefore does not compromise the generality of our design.

We evaluate the design of BASALISC for correctness and performance. Our functional units are emulated and formally verified to meet their specification. We also simulate performance on two FHE benchmarks, showing more than 4,000 times speedup compared to classical software implementations. In future work, we aim to put these results into practice via fabrication of an IC that can be applied in real-world applications.

ACKNOWLEDGMENTS

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR0011-21-C-0034. The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. This work was additionally supported in part by CyberSecurity Research Flanders with reference number VR20192203 and the Research Council KU Leuven (C16/15/058).

Michiel Van Beirendonck is funded by Research Foundation – Flanders (FWO) as Strategic Basic (SB) PhD fellow (project number 1SD5621N).

REFERENCES

- [1] 2022. Lattigo v3. Online: <https://github.com/tuneinsight/lattigo>. (April 2022). EPFL-LDS, Tune Insight SA.
- [2] Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman. 1974. *The Design and Analysis of Computer Algorithms*. Addison-Wesley.
- [3] Martin Albrecht, Melissa Chase, Hao Chen, Jintai Ding, Shafi Goldwasser, Sergey Gorbunov, Shai Halevi, Jeffrey Hoffstein, Kim Laine, Kristin Lauter, Satya Lokam, Daniele Micciancio, Dustin Moody, Travis Morrison, Amit Sahai, and Vinod Vaikuntanathan. 2018. *Homomorphic Encryption Security Standard*. Technical Report. HomomorphicEncryption.org, Toronto, Canada.
- [4] Jean Claude Bajard and Sylvain Duquesne. 2021. Montgomery-friendly primes and applications to cryptography. *Journal of Cryptographic Engineering* 11, 4 (2021), 399–415.
- [5] Jean-Claude Bajard, Julien Eynard, M Anwar Hasan, and Vincent Zucca. 2016. A full RNS variant of FV like somewhat homomorphic encryption schemes. In *International Conference on Selected Areas in Cryptography*. Springer, 423–442.
- [6] Charlotte Bonte, Iliia Iliashenko, Jeongeun Park, Hilder V. L. Pereira, and Nigel P. Smart. 2022. FINAL: Faster FHE instantiated with NTRU and LWE. *Cryptology ePrint Archive*, Report 2022/074. (2022). <https://ia.cr/2022/074>.
- [7] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. 2014. (Leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)* 6, 3 (2014), 1–36.
- [8] Kyle Carter, Adam Foltzer, Joe Hendrix, Brian Huffman, and Aaron Tomb. 2013. SAW: The Software Analysis Workbench. In *Proceedings of the 2013 ACM SIGAda Annual Conference on High Integrity Language Technology (HILT '13)*. Association for Computing Machinery, New York, NY, USA, 15–18. <https://doi.org/10.1145/2527269.2527277>
- [9] Donald Donglong Chen, Nele Mentens, Frederik Vercauteren, Sujoy Sinha Roy, Ray C. C. Cheung, Derek Pao, and Ingrid Verbauwhede. 2015. High-Speed Polynomial Multiplication Architecture for Ring-LWE and SHE Cryptosystems. *IEEE Transactions on Circuits and Systems I: Regular Papers* 62, 1 (Jan. 2015), 157–166. <https://doi.org/10.1109/TCSI.2014.2350431> 100 citations (Semantic Scholar/DOI) [2022-04-29].
- [10] Hao Chen, Ran Gilad-Bachrach, Kyoohyung Han, Zhicong Huang, Amir Jalali, Kim Laine, and Kristin Lauter. 2018. Logistic regression over encrypted data from fully homomorphic encryption. *BMC medical genomics* 11, 4 (2018), 3–12.
- [11] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. 2017. Homomorphic Encryption for Arithmetic of Approximate Numbers. In *Advances in Cryptology – ASIACRYPT 2017*, Tsuyoshi Takagi and Thomas Peyrin (Eds.). Springer International Publishing, Cham, 409–437.
- [12] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. 2020. TFHE: Fast Fully Homomorphic Encryption Over the Torus. 33, 1 (Jan. 2020), 34–91. <https://doi.org/10.1007/s00145-019-09319-x>
- [13] Ilaria Chillotti, Marc Joye, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. 2020. CONCRETE: Concrete Operates on Ciphertexts Rapidly by Extending TtHE. In *WAHC 2020–8th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, Vol. 15.
- [14] D. Cohen. 1976. Simplified control of FFT hardware. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 24, 6 (1976), 577–579. <https://doi.org/10.1109/TASSP.1976.1162854>
- [15] Leo de Castro, Rashmi Agrawal, Rabia Yazicigil, Anantha Chandrakanan, Vinod Vaikuntanathan, Chiraag Juvekar, and Ajay Joshi. 2021. Does Fully Homomorphic Encryption Need Compute Acceleration? *arXiv:2112.06396 [cs]* (Dec. 2021). <http://arxiv.org/abs/2112.06396> 0 citations (Semantic Scholar/arXiv) [2022-04-29] arXiv: 2112.06396.
- [16] Yarkan Doröz, Erdinç Öztürk, and Berk Sunar. 2015. Accelerating Fully Homomorphic Encryption in Hardware. *IEEE Trans. Comput.* 64, 6 (June 2015), 1509–1521. <https://doi.org/10.1109/TC.2014.2345388> 72 citations (Semantic Scholar/DOI) [2022-04-29].
- [17] Mario Garrido. 2016. A New Representation of FFT Algorithms Using Triangular Matrices. *IEEE Transactions on Circuits and Systems I: Regular Papers* 63, 10 (2016), 1737–1745. <https://doi.org/10.1109/TCSI.2016.2587822>
- [18] Craig Gentry. 2009. 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*, Michael Mitzenmacher (Ed.). ACM, 169–178. <https://doi.org/10.1145/1536414.1536440>
- [19] Craig Gentry, Shai Halevi, and Nigel P Smart. 2012. Better bootstrapping in fully homomorphic encryption. In *International Workshop on Public Key Cryptography*. Springer, 1–16.
- [20] Craig Gentry, Shai Halevi, and Nigel P Smart. 2012. Fully homomorphic encryption with polylog overhead. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 465–482.

- [21] Craig Gentry, Shai Halevi, and Nigel P Smart. 2012. Homomorphic evaluation of the AES circuit. In *Annual Cryptology Conference*. Springer, 850–867.
- [22] Shai Halevi and Victor Shoup. 2014. Algorithms in HElib. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I (Lecture Notes in Computer Science)*, Juan A. Garay and Rosario Gennaro (Eds.), Vol. 8616. Springer, 554–571. https://doi.org/10.1007/978-3-662-44371-2_31
- [23] Shai Halevi and Victor Shoup. 2021. Bootstrapping for helib. *Journal of Cryptology* 34, 1 (2021), 1–44.
- [24] L.G. Johnson. 1992. Conflict free memory addressing for dedicated FFT hardware. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing* 39, 5 (1992), 312–316. <https://doi.org/10.1109/82.142032>
- [25] Andrey Kim, Yuriy Polyakov, and Vincent Zucca. 2021. Revisiting Homomorphic Encryption Schemes for Finite Fields. In *Advances in Cryptology - ASIACRYPT 2021*, Mehdi Tibouchi and Huaxiong Wang (Eds.). Springer International Publishing, Cham, 608–639.
- [26] Sangpyo Kim, Wonkyung Jung, Jaiyoung Park, and Jung Ho Ahn. 2020. Accelerating Number Theoretic Transformations for Bootstrappable Homomorphic Encryption on GPUs. In *IEEE International Symposium on Workload Characterization, IISWC 2020, Beijing, China, October 27-30, 2020*. IEEE, 264–275. <https://doi.org/10.1109/IISWC50251.2020.00033>
- [27] Jeffrey R Lewis and Brad Martin. 2003. Cryptol: High assurance, retargetable crypto development and validation. In *IEEE Military Communications Conference, 2003. MILCOM 2003.*, Vol. 2. IEEE, 820–825.
- [28] Kun-Shan Lin, G.A. Frantz, and R. Simar. 1987. The TMS320 family of digital signal processors. *Proc. IEEE* 75, 9 (1987), 1143–1159. <https://doi.org/10.1109/PROC.1987.13867>
- [29] Yutai Ma. 1999. An effective memory addressing scheme for FFT processors. *IEEE Trans. Signal Process.* 47, 3 (1999), 907–911. <https://doi.org/10.1109/78.747802>
- [30] Ahmet Can Mert, Aikata, Sunmin Kwon, Youngsam Shin, Donghoon Yoo, Yongwoo Lee, and Sujoy Sinha Roy. 2022. Medha: Microcoded Hardware Accelerator for computing on Encrypted Data. Cryptology ePrint Archive, Report 2022/480. (2022). <https://ia.cr/2022/480>.
- [31] Ahmet Can Mert, Erdiñç Öztürk, and ErKay Savas. 2019. Design and Implementation of a Fast and Scalable NTT-Based Polynomial Multiplier Architecture. In *22nd Euromicro Conference on Digital System Design, DSD 2019, Kallithea, Greece, August 28-30, 2019*. IEEE, 253–260. <https://doi.org/10.1109/DSD.2019.00045>
- [32] Peter L Montgomery. 1985. Modular multiplication without trial division. *Mathematics of computation* 44, 170 (1985), 519–521.
- [33] Yuriy Polyakov, Kurt Rohloff, and Gerard W Ryan. 2017. Palisade lattice cryptography library user manual. (2017).
- [34] Thomas Pöppelmann, Tobias Oder, and Tim Güneysu. 2015. High-Performance Ideal Lattice-Based Cryptography on 8-Bit ATxmega Microcontrollers. In *Progress in Cryptology - LATINCRYPT 2015 - 4th International Conference on Cryptology and Information Security in Latin America, Guadalajara, Mexico, August 23-26, 2015, Proceedings (Lecture Notes in Computer Science)*, Kristin E. Lauter and Francisco Rodriguez-Henriquez (Eds.), Vol. 9230. Springer, 346–365. https://doi.org/10.1007/978-3-319-22174-8_19
- [35] Thomas Pöppelmann, Michael Naehrig, Andrew Putnam, and Adrian Macias. 2015. Accelerating Homomorphic Evaluation on Reconfigurable Hardware. In *Cryptographic Hardware and Embedded Systems - CHES 2015 (Lecture Notes in Computer Science)*, Tim Güneysu and Helena Handschuh (Eds.). Springer, Berlin, Heidelberg, 143–163. https://doi.org/10.1007/978-3-662-48324-4_8 57 citations (Semantic Scholar/DOI) [2022-04-29].
- [36] Dionysios I. Reisis and Nikolaos Vlassopoulos. 2008. Conflict-Free Parallel Memory Accessing Techniques for FFT Architectures. *IEEE Trans. Circuits Syst. I Regul. Pap.* 55-I, 11 (2008), 3438–3447. <https://doi.org/10.1109/TCSI.2008.924889>
- [37] M. Sadegh Riaz, Kim Laine, Blake Pelton, and Wei Dai. 2020. HEAX: An Architecture for Computing on Encrypted Data. In *ASPLOS '20: Architectural Support for Programming Languages and Operating Systems, Lausanne, Switzerland, March 16-20, 2020*, James R. Larus, Luis Ceze, and Karin Strauss (Eds.). ACM, 1295–1309. <https://doi.org/10.1145/3373376.3378523>
- [38] Stephen Richardson, Dejan Marković, Andrew Danowitz, John Brunhaver, and Mark Horowitz. 2015. Building Conflict-Free FFT Schedules. *IEEE Transactions on Circuits and Systems I: Regular Papers* 62, 4 (2015), 1146–1155. <https://doi.org/10.1109/TCSI.2015.2402935>
- [39] Ronald L Rivest, Len Adleman, Michael L Dertouzos, et al. 1978. On data banks and privacy homomorphisms. *Foundations of secure computation* 4, 11 (1978), 169–180.
- [40] Sujoy Sinha Roy, Ahmet Can Mert, Aikata, Sunmin Kwon, Youngsam Shin, and Donghoon Yoo. 2021. Accelerator for Computing on Encrypted Data. *IACR Cryptol. ePrint Arch.* (2021), 1555. <https://eprint.iacr.org/2021/1555>
- [41] Sujoy Sinha Roy, Frederik Vercauteren, Nele Mentens, Donald Donglong Chen, and Ingrid Verbauwhede. 2014. Compact Ring-LWE Cryptoprocessor. In *Cryptographic Hardware and Embedded Systems - CHES 2014 - 16th International Workshop, Busan, South Korea, September 23-26, 2014. Proceedings (Lecture Notes in Computer Science)*, Lejla Batina and Matthew Robshaw (Eds.), Vol. 8731. Springer, 371–391. https://doi.org/10.1007/978-3-662-44709-3_21
- [42] Nikola Samardzic, Axel Feldmann, Aleksandar Krastev, Srinivas Devadas, Ronald Dreslinski, Christopher Peikert, and Daniel Sanchez. 2021. F1: A Fast and Programmable Accelerator for Fully Homomorphic Encryption. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '21)*. Association for Computing Machinery, New York, NY, USA, 238–252. <https://doi.org/10.1145/3466752.3480070>
- [43] SEAL 2022. Microsoft SEAL (release 4.0). <https://github.com/Microsoft/SEAL>. (March 2022). Microsoft Research, Redmond, WA.
- [44] Sujoy Sinha Roy, Kimmo Järvinen, Frederik Vercauteren, Vassil Dimitrov, and Ingrid Verbauwhede. 2015. Modular Hardware Architecture for Somewhat Homomorphic Function Evaluation. In *Cryptographic Hardware and Embedded Systems - CHES 2015 (Lecture Notes in Computer Science)*, Tim Güneysu and Helena Handschuh (Eds.). Springer, Berlin, Heidelberg, 164–184. https://doi.org/10.1007/978-3-662-48324-4_9 40 citations (Semantic Scholar/DOI) [2022-04-29].
- [45] Sujoy Sinha Roy, Kimmo Järvinen, Jo Vliegen, Frederik Vercauteren, and Ingrid Verbauwhede. 2018. HEPCloud: An FPGA-Based Multicore Processor for FV Somewhat Homomorphic Function Evaluation. *IEEE Trans. Comput.* 67, 11 (Nov. 2018), 1637–1650. <https://doi.org/10.1109/TC.2018.2816640> 25 citations (Semantic Scholar/DOI) [2022-04-29].
- [46] Sujoy Sinha Roy, Furkan Turan, Kimmo Järvinen, Frederik Vercauteren, and Ingrid Verbauwhede. 2019. FPGA-Based High-Performance Parallel Architecture for Homomorphic Computing on Encrypted Data. In *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. 387–398. <https://doi.org/10.1109/HPCA.2019.00052> 50 citations (Semantic Scholar/DOI) [2022-04-29].
- [47] Nigel P Smart and Frederik Vercauteren. 2014. Fully homomorphic SIMD operations. *Designs, codes and cryptography* 71, 1 (2014), 57–81.
- [48] Pei-Yun Tsai and Chung-Yi Lin. 2011. A Generalized Conflict-Free Memory Addressing Scheme for Continuous-Flow Parallel-Processing FFT Processors With Rescheduled. *IEEE Trans. Very Large Scale Integr. Syst.* 19, 12 (2011), 2290–2302. <https://doi.org/10.1109/TVLSI.2010.2077314>
- [49] Furkan Turan, Sujoy Sinha Roy, and Ingrid Verbauwhede. 2020. HEAWS: An Accelerator for Homomorphic Encryption on the Amazon AWS FPGA. *IEEE Trans. Computers* 69, 8 (2020), 1185–1196. <https://doi.org/10.1109/TC.2020.2988765>
- [50] Furkan Turan, Sujoy Sinha Roy, and Ingrid Verbauwhede. 2020. HEAWS: An Accelerator for Homomorphic Encryption on the Amazon AWS FPGA. *IEEE Trans. Comput.* 69, 8 (Aug. 2020), 1185–1196. <https://doi.org/10.1109/TC.2020.2988765> 13 citations (Semantic Scholar/DOI) [2022-04-29].
- [51] Vincent Zucca. 2018. *Towards efficient arithmetic for Ring-LWE based homomorphic encryption*. Ph.D. Dissertation. Sorbonne université.

A KEY SWITCHING PROCEDURE

Key switching transforms a ciphertext ct encrypting a plaintext m under a key s , into a ciphertext ct' encrypting the same plaintext under a different key s' . In practice, we usually switch from s^2 (for multiplication) or from $\phi_k(s)$ (for automorphism) to the original secret key. We need an auxiliary data structure called *key switching matrix*, which is essentially a set of encryptions of s under s' . A high level description of various key switching methods is given in Appendix B by Kim et al. [25]. BASALISC implements the *hybrid key switching* method from Appendix B.2.3, which we review here.

A.1 Notations

Consider two coprime moduli

$$Q = \prod_{i=1}^{\ell} q_i \quad \text{and} \quad P = \prod_{i=\ell+1}^{\ell+k} q_i. \quad (4)$$

The input ciphertext is defined modulo Q , but the key switching matrix will be defined with respect to the extended modulus QP .

A main concept of key switching is *digit decomposition*: we fix some $d \in \mathbb{N}$ (typically between 3 and 5) and define “base digits” as

$$D_i = \prod_{j=(i-1) \cdot (\ell/d) + 1}^{i \cdot \ell/d} q_j,$$

In this appendix, ℓ does not denote the number of slots.

for $1 \leq i \leq d$. For simplicity, we assume that ℓ is divisible by d . Note that each base digit is a product of ℓ/d primes q_j . Also let

$$\hat{D}_i = Q/D_i, \quad \hat{D}_i^{-1} = (Q/D_i)^{-1} \pmod{D_i},$$

and

$$\vec{D} = \left([\hat{D}_1^{-1} \cdot \hat{D}_1]_Q, \dots, [\hat{D}_d^{-1} \cdot \hat{D}_d]_Q \right)^\top.$$

We define the digit decomposition of $\mathbf{u} \in \mathcal{R}_Q$ as

$$D^{-1}(\mathbf{u}) = ([\mathbf{u}]_{D_1}, \dots, [\mathbf{u}]_{D_d})^\top.$$

An important observation is that

$$\langle D^{-1}(\mathbf{u}), \vec{D} \rangle = \mathbf{u} \pmod{Q},$$

where $\langle \cdot, \cdot \rangle$ denotes the dot product between left and right vector.

A.2 Fast Base Extension

During key switching, the input ciphertext is defined modulo Q , but the key switching matrix is defined modulo QP . Therefore, we need a procedure that extends the modulus temporarily by including all factors of P . This procedure is called *fast base extension*, and it is shown in Algorithm 1. Given a polynomial $\mathbf{a} \in \mathcal{R}_Q$ in coefficient representation, we extend it from \mathcal{R}_Q to \mathcal{R}_{QP} as follows:

$$\text{FastBaseExt}(\mathbf{a}, Q, P) = \left(\left[\sum_{i=1}^{\ell} \left[\mathbf{a} \cdot \left(\frac{Q}{q_i} \right)^{-1} \right]_{q_i} \cdot \frac{Q}{q_i} \right]_{q_j} \right)_{j=\ell+1}^{\ell+k}.$$

The algorithm assumes that

$$Q_i = \frac{Q}{q_i} \quad \text{and} \quad Q_i^{-1} = \left(\frac{Q}{q_i} \right)^{-1} \pmod{q_i}$$

are given in precomputed format. Note that the residue of \mathbf{a} moduli q_i is denoted by $\mathbf{a}^{(i)}$, and that the accumulator on line 5 is implicitly initialized to 0.

Algorithm 1 Fast base extension

Require: $\mathbf{a} \in \mathcal{R}_Q, Q$ and P

Ensure: $\mathbf{a} \in \mathcal{R}_{QP}$

```

1: function FASTBASEEXT( $\mathbf{a}, Q, P$ )
2:   for  $i \in \{1, \dots, \ell\}$  do
3:      $\mathbf{r} \leftarrow [\mathbf{a}^{(i)} \cdot Q_i^{-1}]_{q_i}$ 
4:     for  $j \in \{\ell+1, \dots, \ell+k\}$  do
5:        $\mathbf{a}^{(j)} \leftarrow [\mathbf{a}^{(j)} + \mathbf{r} \cdot Q_i]_{q_j}$        $\triangleright$  MAC unit
6:     end for
7:   end for
8:   return  $(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(\ell+k)})$ .
9: end function

```

A.3 Hybrid Key Switching

The high level idea of hybrid key switching is to provide a key switching matrix that encrypts $\vec{D} \cdot \mathbf{s}$ under \mathbf{s}' . Now given a ciphertext

$$\text{ct} = (c_0, \mathbf{c}_1) = (-c_1 \cdot \mathbf{s} + t\mathbf{e} + \mathbf{m}, \mathbf{c}_1) \in \mathcal{R}_{QP}^2,$$

we can decompose \mathbf{c}_1 into digits and multiply by the key switching matrix to obtain an encryption of $\mathbf{c}_1 \cdot \mathbf{s}$ under \mathbf{s}' . We add the result to c_0 to remove the term $-c_1 \cdot \mathbf{s}$, and obtain $\mathbf{c}'_0 = -\mathbf{c}'_1 \cdot \mathbf{s}' + t\mathbf{e}' + \mathbf{m}$.

Following the above analysis, the key switching matrix is generated by the client as

$$\vec{\text{evk}} = (\vec{\mathbf{k}}_0, \vec{\mathbf{k}}_1) \in \mathcal{R}_{QP}^{d \times 2},$$

with $\vec{\mathbf{k}}_0 = -\vec{\mathbf{k}}_1 \cdot \mathbf{s}' + t\vec{\mathbf{e}} + P\vec{D} \cdot \mathbf{s}$. To key switch a ciphertext, we compute

$$\vec{\mathbf{y}} = D^{-1}(\mathbf{c}_1) \in \mathcal{R}_Q^d,$$

and multiply it by both columns of the matrix $\vec{\text{evk}}$. Notice that $\vec{\mathbf{y}}$ is defined modulo Q , but the multiplications are performed with respect to QP . This is where the fast base extension from Algorithm 1 is used as a subroutine. So finally, key switching computes

$$\tilde{\text{ct}} = (\tilde{c}_0, \tilde{\mathbf{c}}_1) = \left(\langle \vec{\mathbf{y}}, \vec{\mathbf{k}}_0 \rangle, \langle \vec{\mathbf{y}}, \vec{\mathbf{k}}_1 \rangle \right) \in \mathcal{R}_{QP}^2, \quad (5)$$

and adds it to the original ciphertext as

$$\text{ct}' = \left(\left[c_0 + \frac{\tilde{c}_0 + t\delta_0}{P} \right]_Q, \left[\tilde{\mathbf{c}}_1 + \frac{t\delta_1}{P} \right]_Q \right) \in \mathcal{R}_Q^2, \quad (6)$$

with $\delta_i = [-t^{-1}\tilde{c}_i]_P$. Note that the division by P should be interpreted as a modulus switching step, which is necessary to bring the ciphertext modulus back to Q .

Key switching relies on the NTT unit to convert between coefficient format and Double-CRT. In particular, fast base extension can only be done in coefficient format, and is necessary in Equation 5 (to raise the ciphertext from Q to QP) and in Equation 6 (to raise δ_i from P to QP).

B BOOTSTRAPPING DETAILS

This appendix includes extra details about our Montgomery-friendly bootstrapping algorithm. We first give the proof of Lemma 3.1, presented in section 3.3. Then we give pseudocode for our new bootstrapping technique. Finally, we explain how the parameters of our method can be chosen in practice.

B.1 Proof of Lemma 3.1

LEMMA 3.1. *Let $p > 1$ be a prime number, and let $e > r \geq 1$ and $q = 1 \pmod{p^e}$ be sufficiently high parameters. If (c_0, \mathbf{c}_1) is a BGV encryption of \mathbf{m} with plaintext modulus p^r and ciphertext modulus q , then we can decrypt it by computing*

$$\mathbf{c}'_i \leftarrow [p^{e-r} \mathbf{c}_i]_q, \quad \mathbf{w} \leftarrow [\mathbf{c}'_0 + \mathbf{c}'_1 \cdot \mathbf{s}]_{p^e} \quad \text{and} \quad \mathbf{m} \leftarrow [[\mathbf{w}/p^{e-r}]_{p^r}].$$

Here we use $\lfloor \cdot \rfloor$ for coefficient-wise rounding to the nearest integer.

PROOF. Let $\mathbf{u} = \mathbf{c}'_0 + \mathbf{c}'_1 \cdot \mathbf{s}$, then it follows that

$$\mathbf{u} = p^{e-r}(\mathbf{c}_0 + \mathbf{c}_1 \cdot \mathbf{s}) = p^{e-r}(\mathbf{m} + p^r \mathbf{e}) \pmod{q},$$

where we have used the definition of \mathbf{c}'_i and Equation 1. Now we make the reduction modulo q explicit and write

$$\mathbf{u} = p^{e-r}(\mathbf{m} + p^r \mathbf{e}) + q\mathbf{r} \quad (7)$$

for some $\mathbf{r} \in \mathcal{R}$. Following the decryption procedure, we have

$$\mathbf{w} = [\mathbf{u}]_{p^e} = [p^{e-r}(\mathbf{m} + p^r \mathbf{e}) + q\mathbf{r}]_{p^e} = [p^{e-r}\mathbf{m} + \mathbf{r}]_{p^e},$$

where we have used $q = 1 \pmod{p^e}$. Now we make the reduction modulo p^e explicit and write

$$\mathbf{w} = p^{e-r}\mathbf{m} + \mathbf{r} + p^e \mathbf{t}$$

for some $t \in \mathcal{R}$. Again following the decryption procedure, we have

$$[\lfloor \mathbf{w}/p^{e-r} \rfloor]_{p^r} = [\mathbf{m} + \lfloor \mathbf{r}/p^{e-r} \rfloor]_{p^r} = \mathbf{m}$$

where the last equation is correct if the coefficients of \mathbf{r} are upper bounded by $p^{e-r}/2$. Formally, we write this requirement as $\|\mathbf{r}\|_\infty \leq p^{e-r}/2$, where $\|\mathbf{r}\|_\infty$ denotes the uniform norm on the coefficients of \mathbf{r} . So we need to find parameters e and q that satisfy this requirement.

Applying the triangle inequality on Equation 7, we have

$$\begin{aligned} \|\mathbf{r}\|_\infty &\leq \|\mathbf{u}/q\|_\infty + \|p^{e-r}(\mathbf{m} + p^r \mathbf{e})/q\|_\infty \\ &\leq \|(\mathbf{c}'_0 + \mathbf{c}'_1 \cdot \mathbf{s})/q\|_\infty + \|p^{e-r} \mathbf{m}/q\|_\infty + \|p^e \mathbf{e}/q\|_\infty. \end{aligned} \quad (8)$$

The third term on the right-hand side of Equation 8 depends on the remaining noise budget of $(\mathbf{c}_0, \mathbf{c}_1)$. Formally, we impose that the noise rate is upper bounded as $\|\mathbf{e}/q\|_\infty \leq B_1$. The second term can be made arbitrarily small by taking q sufficiently high. Formally, we impose $\|\mathbf{m}/q\|_\infty \leq B_2$. For the first term, recall that the coefficients of \mathbf{s} are small, so let them be upper bounded by B_3 . Also note that \mathbf{c}'_0 and \mathbf{c}'_1 have coefficients in the set $[-q/2, q/2] \cap \mathbb{Z}$. Moreover, it is a well-known fact that multiplication modulo $X^N + 1$ cannot increase the norm more than a factor of N (e.g., see Zucca [51]). We combine these three insights to get

$$\|(\mathbf{c}'_0 + \mathbf{c}'_1 \cdot \mathbf{s})/q\|_\infty \leq (1 + B_3 \cdot N)/2.$$

Combining the three upper bounds, Equation 8 reduces to

$$\|\mathbf{r}\|_\infty \leq (1 + B_3 \cdot N)/2 + p^{e-r} B_2 + p^e B_1.$$

Recall that our requirement was $\|\mathbf{r}\|_\infty \leq p^{e-r}/2$, so it suffices to find parameters e and q such that

$$(1 + B_3 \cdot N)/2 + p^{e-r} B_2 + p^e B_1 \leq p^{e-r}/2. \quad (9)$$

A set of example parameters shows that this inequality can be satisfied: take $B_1 = p^{-r}/8$ and $B_2 = 1/8$, then we are left with

$$(1 + B_3 \cdot N)/2 \leq p^{e-r}/4.$$

Finally, we can choose the smallest possible value of e that satisfies this inequality and choose q accordingly. \square

B.2 Pseudocode

Our Montgomery-friendly bootstrapping needs one more subroutine that is known as small Montgomery reduction. It was introduced by Bajard et al. [5] and repeated here in Algorithm 2. It takes as input a polynomial $\mathbf{a} \in \mathcal{R}_{QP}$ in coefficient representation, and outputs $\mathbf{a} \cdot P^{-1} \in \mathcal{R}_Q$ with coefficients reduced modulo a given parameter m . Specifically, the coefficients will be upper bounded by $(1 + \epsilon)m/2$ for some $\epsilon \ll 1$ that is not further specified here. We need this subroutine for reduction modulo q and p^e in Lemma 3.1. Note that the algorithm is defined with respect to Q and P as in Equation 4, and that the residue of \mathbf{a} moduli q_i is denoted by $\mathbf{a}^{(i)}$.

Algorithm 3 gives the pseudocode for our Montgomery-friendly bootstrapping. This is a direct translation of Lemma 3.1 to the homomorphic domain, and it includes the following steps:

This small Montgomery reduction is not directly related to the fact that we use Montgomery multipliers. In fact, we have even specified Algorithm 1 and Algorithm 2 assuming a standard reduction technique. When instantiating either algorithm with Montgomery multipliers, we need to convert all residues out and in Montgomery format whenever we reinterpret a variable modulo q_i as a variable modulo q_j . In both algorithms, this happens on the fifth line.

Algorithm 2 Small Montgomery reduction

Require: $\mathbf{a} \in \mathcal{R}_{QP}$, Q , P and m s.t. $\|\mathbf{a}\|_\infty \ll P \cdot m$

Ensure: $\mathbf{b} \in \mathcal{R}_Q$ s.t. $\mathbf{b} = \mathbf{a} \cdot P^{-1} \pmod{m}$ and $\|\mathbf{b}\|_\infty \leq (1 + \epsilon)m/2$

```

1: function SMALLMONTRED( $\mathbf{a}$ ,  $Q$ ,  $P$ ,  $m$ )
2:   for  $i \in \{\ell + k, \dots, \ell + 1\}$  do       $\triangleright$  Loop in reverse direction
3:      $\mathbf{r} \leftarrow [-\mathbf{a}^{(i)} \cdot m^{-1}]_{q_i}$ 
4:     for  $j \in \{1, \dots, i - 1\}$  do
5:        $\mathbf{a}^{(j)} \leftarrow [(\mathbf{a}^{(j)} + m \cdot \mathbf{r}) \cdot q_i^{-1}]_{q_j}$        $\triangleright$  MAC unit
6:     end for
7:   end for
8:   return  $(\mathbf{a}^{(1)}, \dots, \mathbf{a}^{(\ell)})$ 
9: end function

```

- Line 4 multiplies the input ciphertext by p^{e-r} and an auxiliary modulus b . The auxiliary modulus is introduced for compensation on line 7.
- Line 6 extends the ciphertext modulus from q to $Q \cdot q \cdot b$ using fast base extension from Algorithm 1. This procedure can lead to undesired overflows modulo q , which causes $\|\mathbf{d}_i\|_\infty$ to be greater than $q/2$. This increases the tightness of the bound in Equation 8, but fortunately, the overflows can be compensated in the next step. Finally, note that this step assumes that Q , q and b are pairwise coprime.
- Line 7 compensates for possible overflows modulo q introduced on line 6. The small Montgomery reduction has two side effects: it decreases the modulus from $Q \cdot q \cdot b$ to $Q \cdot q$, and the result gets an additional factor $b^{-1} \pmod{q}$. The latter was already compensated by the factor b on line 4.
- Line 8 reduces the result modulo p^e and further decreases the modulus from $Q \cdot q$ to q . Now the result gets no additional factor since $q = 1 \pmod{p^e}$. Note that this step is necessary to minimize the noise growth in the next step.
- Line 10 takes the inner product between the ciphertext and the secret key. The secret key is processed homomorphically in the form of a bootstrapping key.
- Line 11 performs homomorphic coefficient-wise rounding. This functionality is the same as in HELib, so we can reuse its implementation. Note that this step dominates execution time in practice.

B.3 Choice of Parameters

We have a few notes on the concrete choice of e and q :

- The constraint from Equation 9 is in practice determined by the first term. The reason is that we can choose B_1 and B_2 significantly lower than B_3 , by taking q sufficiently high and preventing the noise from growing to its maximum level. Hence the concrete values of e and q mainly depend on the secret key distribution and the ring dimension N .
- The parameter selection from the proof is rigorous, so decryption succeeds with 100% probability. However, the complexity of bootstrapping increases with the magnitude of e , so it is beneficial to take it as low as possible. Halevi and Shoup [23] have therefore proposed a statistical analysis on the first term of Equation 8. Leveraging their approach, we

Algorithm 3 Montgomery-friendly bootstrapping

Require: $ct = (c_0, c_1) \in \mathcal{R}_Q^2$ and $bsk \in \mathcal{R}_Q^2$ s.t. $q = 1 \pmod{p^e}$

Ensure: $ct' \in \mathcal{R}_Q^2$ s.t. $\text{Dec}(ct') = \text{Dec}(ct)$

```

1: function BOOTSTRAP(ct, bsk)
2:   for  $i \in \{0, 1\}$  do
3:     for  $j \in \{\ell + 1, \dots, \ell + k\}$  do
4:        $d_i^{(j)} \leftarrow [c_i^{(j)} \cdot p^{e-r} \cdot b]_{q_j}$ 
5:     end for
6:      $d_i \leftarrow \text{FASTBASEEXT}(d_i, q, Q \cdot b)$ 
7:      $c'_i \leftarrow \text{SMALLMONTRED}(d_i, Q, q, b, q)$   $\triangleright$  Reduce mod  $q$ 
8:      $c'_i \leftarrow \text{SMALLMONTRED}(c'_i, Q, q, p^e)$   $\triangleright$  Reduce mod  $p^e$ 
9:   end for
10:   $ct' \leftarrow \text{ADD}(\text{MUL}(bsk, c'_1), c'_0)$ 
11:  return  $\lfloor ct'/p^{e-r} \rfloor$   $\triangleright$  Same as in HELib
12: end function
    
```

can choose e based on a trade-off between time complexity and probability of a bootstrapping failure.

- For a plaintext modulus of 15 bits or less, we can directly take q as a Montgomery-friendly prime that satisfies the constraint of Lemma 3.1. For higher precision plaintext spaces, this is not directly possible anymore since the native word size of BASALISC is 32 bits, and the constraint $q = 1 \pmod{2N}$ already consumes 17 bits. Hence we must apply a brute force or meet-in-the-middle search for an appropriate q that factors into 32-bit Montgomery-friendly

primes. Since this is a tedious procedure, it would be beneficial to weaken the constraint of Lemma 3.1 to $\gcd(q, p^e) = 1$. This is possible, provided that we change the last equation in Lemma 3.1 to

$$m \leftarrow [q \cdot \lfloor q^{-1} \cdot w/p^{e-r} \rfloor]_{p^r}$$

for $q^{-1} \cdot q = 1 \pmod{p^e}$. We omit the adapted proof.

C COMPARISON TO HELIB

Table 8 compares BASALISC performance to HELib for the NTT, and some basic and auxiliary homomorphic operations. Each operand is a freshly encrypted ciphertext using the example parameter set from Table 1. Note that the NTT benchmark converts an entire ciphertext from coefficient representation to Double-CRT.

We achieve major speedups for all homomorphic operations. In particular, we accelerate key switching - the most time-intensive kernel - by a factor of $2.0 \cdot 10^3$.

Table 8: Performance comparison of HELib and BASALISC.

Operation	HELlib	BASALISC	Speedup
NTT	27 ms	11 μ s	$2.5 \cdot 10^3 \times$
Add/Sub	4 ms	8 μ s	$5.0 \cdot 10^2 \times$
Plaintext mul	159 ms	5 μ s	$3.2 \cdot 10^4 \times$
Mul (no key switch)	531 ms	20 μ s	$2.7 \cdot 10^4 \times$
Permutation (no key switch)	12 ms	11 μ s	$1.1 \cdot 10^3 \times$
Key switching	580 ms	292 μ s	$2.0 \cdot 10^3 \times$