

CARM: CUDA-Accelerated RNS Multiplication in Word-Wise Homomorphic Encryption Schemes

Shiyu Shen¹, Hao Yang², Yu Liu¹, Zhe Liu², and Yunlei Zhao¹

¹ School of Computer Science, Fudan University, Shanghai, China
{shenshiyu21,yu_liu21}@m.fudan.edu.cn, ylzhao@fudan.edu.cn

² College of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, Nanjing, China
crypto@d4rk.dev, zhe.liu@nuaa.edu.cn

Abstract. Homomorphic encryption (HE), which allows computation over encrypted data, has often been used to preserve privacy. However, the computationally heavy nature and complexity of network topologies make the deployment of HE schemes in the Internet of Things (IoT) scenario difficult. In this work, we propose CARM, the first optimized GPU implementation that covers BGV, BFV and CKKS, targeting for accelerating homomorphic multiplication using GPU in heterogeneous IoT systems. We offer constant-time low-level arithmetic with minimum instructions and memory usage, as well as performance- and memory-prior configurations, and exploit a parametric and generic design, and offer various trade-offs between resource and efficiency, yielding a solution suitable for accelerating RNS homomorphic multiplication on both high-performance and embedded GPUs. Through this, we can offer more real-time evaluation results and relieve the computational pressure on cloud devices. We deploy our implementations on two GPUs and achieve up to 378.4×, 234.5×, and 287.2× speedup for homomorphic multiplication of BGV, BFV, and CKKS on Tesla V100S, and 8.8×, 9.2×, and 10.3× on Jetson AGX Xavier, respectively.

Keywords: Homomorphic encryption · RNS multiplication · Number Theoretic Transform · Internet of Things · GPU acceleration.

1 Introduction

The Internet of Things (IoT) is now infiltrating into all parts of people’s lives. The resource-constrained nature of IoT edge devices makes it necessary for cloud servers to assist in processing data in some scenarios, which leads to security issues when the data is privacy-sensitive. In detail, processing user data in plaintext form will leak personal information that could be used to track user interests, locations, and so on. Homomorphic encryption (HE) is often introduced to solve this problem, which allows data analysis to be performed on ciphertext and enables secure device-to-cloud computation.

Since Gentry proposed the first fully homomorphic encryption (FHE) scheme [27, 28], there have been many improvements to make FHE more practical, such as enabling Single Instruction Multiple Data (SIMD) operations [42] and handling floating point data [19]. In view of the good parallelism of the low-level arithmetic of HE schemes, parallel computing is usually used in practice to speed up the computing process, relieving the computational overhead of the schemes. The Graphics Processing Unit (GPU) has powerful parallel computing capability. Existing approaches that take full advantage of GPUs to accelerate compute-intensive computation achieve auspicious performance, underlining their tremendous impact.

However, the computation overhead is a significant consideration that limits the deployability of homomorphic encryption, especially on resource-constrained IoT edge devices. Additionally, the complexity of IoT topologies and the diversity of the computational power of devices make the implementation more difficult. Fortunately, the massive increase in the quantities and varieties of devices has promoted the evolution of IoT topologies, making the modern IoT systems become more heterogeneous. The emergence of devices with embedded Graphics Processing Units (GPUs) brings new challenges and opportunities to the deployment of homomorphic encryption in IoT scenarios. Embedded GPUs are a promising technology, as they provide high performance computing platform for massive data processing with low energy consumption, while preserving

the portable and mobile nature of IoT devices. Enabling homomorphic encryption on such devices allows computation in the entire device-to-cloud scenario to be accelerated using GPUs, which is of great value for the deployment of FHE schemes in the IoT systems.

1.1 Related work

In 2009, Gentry proposed the first FHE scheme based on ideals [27, 28]. Since then, research in this area began to burgeon, mainly focusing on improving the efficiency and applicability of FHE schemes. Modern FHE schemes are based on different hardness assumptions, i.e., Approximate Greatest Common Divisor (AGCD) [36], (Ring-)Learning with Errors ((R-)LWE) [17, 18], and Number Theory Research Unit (NTRU) [23], and most of them can be classified into bit-wise FHE (e.g., FHEW [25] and TFHE [20]) and word-wise FHE (e.g., BGV [16], BFV [26] and CKKS [19]). Compared to bit-wise FHE, word-wise FHE is more efficient in large vectorial arithmetic operations [15]. Recently, many open source libraries have implemented these schemes, such as HELib [2], PALISADE [3], SEAL [39] and HEAAN [1].

To alleviate the computation bottlenecks, some studies dedicated to exploring different acceleration strategies for polynomial arithmetic, such as adopting the Number Theoretic Transformation (NTT) and the Residue Number System (RNS) representation. SEAL employs the Harvey’s algorithm [32] and follows the classic level-by-level method for implementing NTT, which is further accelerated in several studies [4, 38] using GPUs. Recent studies [9, 22, 33] adopted the four-step Cooley-Tukey algorithm [10, 21] and split the N -point NTT into several ones with smaller size, they also achieved performance improvements on GPUs. The RNS is a commonly utilized technique to efficiently conduct computation over integers larger than the processor word-size. In practice, there are some RNS variant of these schemes, using the RNS instantiations such as the Bajard-Eynard-Hasan-Zucca (BEHZ) method [11] that is based on the decomposition to residues, the Gentry-Halevi-Smart (GHS) method [29] that temporarily extends the basis with a special prime, and the hybrid method [31] that takes advantage of both methods.

Previous works on accelerating FHE schemes mostly focus on high-performance platforms, by increasing the parallelism of computation. For example, Boemer et al. proposed Intel HEXL [13] that accelerated the polynomial arithmetic using the Intel Advanced Vector Extensions 512 (AVX512) instruction set and achieved a $7.2\times$ speedup, which is now integrated into SEAL and PALISADE. Other works dedicated to improving the computation concurrency through general-purpose GPU (GPGPU) implementation, resulting in speedup such as $5\times$ - $22\times$ for homomorphic operations in [4] and $257\times$ for CKKS bootstrapping in [33]. FHE schemes were not introduced to the IoT systems, until Natarajan et al. proposed the first CKKS implementation in the embedded domain, i.e., SEAL-Embedded [37]. They concentrated on reducing the memory consumption of CKKS encoding and encryption while maintaining the performance and used an adapter as a middle layer to convert the ciphertexts sent by the devices into a form compatible to SEAL.

So far, although embedded GPUs have been deployed in many IoT devices, there is no optimized implementation of FHE schemes targeted for this platform. Previous works on accelerating FHE schemes using GPU only support a portion of the parameters and are not general enough to cover all applications. Additionally, many open-source libraries and works adopt non-constant-time implementations, such as SEAL [39] and [33]. A recent side-channel attack on SEAL v3.2 [7], although not affecting implementations after v3.3, raised awareness of the potential security risks that non-constant-time implementations may bring.

1.2 Contributions and road-map

In this work, we propose CARM, the first optimized GPU implementation of RLWE-based FHE schemes in IoT scenario that covers BGV, BFV and CKKS, and integrate it into SEAL v4.0 library [39]. Our library provides fast homomorphic multiplication for two types of devices in the IoT network topology, i.e., middle-layer and cloud devices. Our solution has two advantages: first, we accelerate the performance of homomorphic multiplication on cloud devices, and second, we allow middle-layer devices to provide more real-time evaluation results, thus relieving the computational pressure on cloud devices. We follow the secure deployment solution for IoT [37], and the computation model is shown in Figure 1. In summary, we make the following contributions:

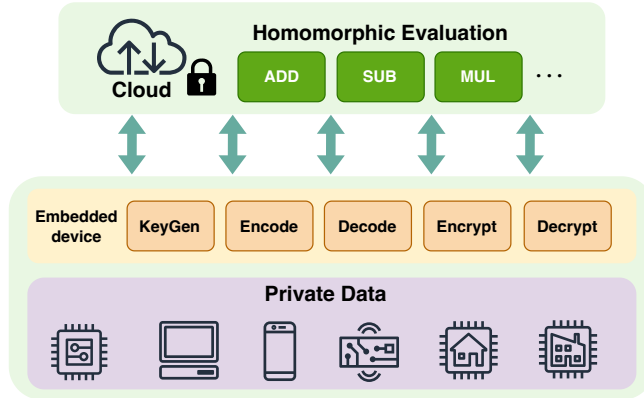


Fig. 1: Secure device-to-cloud computation model of CARM.

- We provide constant-time low-level arithmetic implementation written in the CUDA PTX assembly with minimum instructions and register usage, covering multiply-accumulation, conditional subtraction, Barrett reduction, and butterfly operation, which allow us to eliminate the divergence in warps and provide side-channel resistance.
- We exploit a parametric hierarchical design for the (I)NTT with a mix-type memory usage strategy to fully utilize the on-board memory and reduce the IO latency. We offer the implementation of both performance-prior (PP) and memory-prior (MP) versions, allowing a flexible choice according to the target platforms.
- We devise a generic design of homomorphic multiplication and apply memory-centric optimization through kernel fusing (KF) and data re-use, which achieve a balance between the instruction- and thread-level parallelism. For that KF will not always improve the performance, we illustrate its impact and give the bound between the methods with or without KF. We also exploit the KF approach in the BEHZ-type multiplication that reduces the memory access and meanwhile preserves good parallelism.
- We deploy our GPU implementation on two devices, targeting two different scenarios. Our implementation outperforms the recent published works [4, 6, 8, 33, 35]. Meanwhile, compared with the CPU baseline, our (I)NTT implementation improves the performance up to 140.8× for a single (I)NTT and 251.4× for batched (I)NTT with size 21 on a Tesla V100S GPU, and up to 5.7× on a Jetson AGX Xavier GPU. With the MP configuration, we show a 170.7× improvement in the size of the precomputed table. We also achieve up to 378.4×, 233.6×, and 287.2× on Tesla V100S (8.8×, 9.2×, and 10.3× on Jetson AGX Xavier) for the multiplication of BGV, BFV and CKKS, respectively.

The rest of this paper is organized as follows. In Section 2, we present some definitions and introductions to RLWE-based FHE schemes, embedded GPU, and the programming model. In Section 3, we describe the structure of CARM as well as the considerations and trade-offs in our design. The implementation and optimization details are presented in Section 4. Thereafter, we provide the performance results, comparison, and some further discussions in Section 5. Finally, we conclude this paper in Section 6.

2 Preliminaries

2.1 Notation

Let N be a power of 2. We choose the $2N$ -th cyclotomic polynomial $x^N + 1$ to define the quotient ring $R = \mathbb{Z}[x]/(x^N + 1)$, whose elements are integer polynomials of degree less than N and denoted by lower-case boldface letters, e.g., \mathbf{p} or $p(x)$. For a real number r , we use $\lfloor r \rfloor$, $\lceil r \rceil$, and $\lceil r \rceil$ to denote rounding down, up, and to the nearest integer, respectively, and use $\lceil r \rceil_q$ (for q an integer) to denote the centered remainder of r

modulo q . When these operations are applied to a polynomial, we indicate that the corresponding operation is performed on each coefficient separately. By $a \stackrel{\$}{\leftarrow} S$, we denote that a is sampled uniformly from a finite set S . For a distribution \mathcal{X} on S , we use $a \leftarrow \mathcal{X}$ to denote the sampling of a from S according to the distribution \mathcal{X} .

2.2 BGV, BFV and CKKS

The BGV, BFV and CKKS are three prominent word-wise FHE schemes. This type of scheme supports batched computation by pre-splitting the input into different slots and is efficient on some time-consuming operations. The BGV and BFV work on the same plaintext space that perform exact computations over finite fields, while the CKKS supports approximate evaluation of real numbers at a preset precision. The security of BGV, BFV and CKKS is based on the hardness of the (Ring-)Learning with Errors ((R-)LWE) problem. Let λ denote the security parameter. Let \mathbf{s} be a random element in R_q and $\mathcal{X} = \mathcal{X}(\lambda)$ be a distribution over R_q . The decision R-LWE problem is defined as to distinguish the samples $(\mathbf{a}_i, \mathbf{b}_i)$ from the uniform distribution on R_q^2 , where $\mathbf{b}_i = \mathbf{a}_i \mathbf{s} + \mathbf{e}_i$, $\mathbf{a}_i \stackrel{\$}{\leftarrow} R_q$, $\mathbf{e}_i \leftarrow \mathcal{X}$.

Let $Q = \prod_{i=0}^L q_i$ be the moduli and define the modulus at level l as $Q_l = \prod_i^l q_i$, where q_i are primes, and let p be the special prime [29]. We define the main components of the three algorithms as follows.

- Key generation. Given the system parameter $\mathbf{params} = (\lambda, N, Q)$ and distributions \mathcal{X}_{key} and \mathcal{X}_{err} , generates the public key, secret key and evaluation key in the following way:
 - Secret key. Sample $\mathbf{s} \leftarrow \mathcal{X}_{\text{key}}$, and set the secret key $\mathbf{sk} := (1, \mathbf{s})$.
 - Public key. Sample $\mathbf{a} \stackrel{\$}{\leftarrow} R_Q$ and $\mathbf{e} \leftarrow \mathcal{X}_{\text{err}}$. The public key is formed as $\mathbf{pk} := ([-\mathbf{a} \cdot \mathbf{s} + \mathbf{t}\mathbf{e}]_Q, \mathbf{a})$ in BGV, and as $\mathbf{pk} := ([-\mathbf{a} \cdot \mathbf{s} + \mathbf{e}]_Q, \mathbf{a})$ in both BFV and CKKS.
 - Evaluation key. Sample $\mathbf{a}' \stackrel{\$}{\leftarrow} R_{p \cdot Q}$ and $\mathbf{e}' \leftarrow \mathcal{X}_{\text{err}}$. The evaluation key is defined as $\mathbf{evk} := ([-\mathbf{a}' \cdot \mathbf{s} + \mathbf{e}' + p\mathbf{s}']_{p \cdot Q}, \mathbf{a}')$, where $\mathbf{s}' = [\mathbf{s}^2]_Q$.
- Encryption. Given a public key $\mathbf{pk} = (\mathbf{u}_0, \mathbf{u}_1) \in R_Q^2$ and a message $m \in R$, sample $\mathbf{r} \leftarrow \mathcal{X}_{\text{key}}$ and $\mathbf{e}_0, \mathbf{e}_1 \leftarrow \mathcal{X}_{\text{err}}$. The resulting ciphertext of BGV, BFV, or CKKS encryption is described respectively as follows:

BGV. $\text{Enc}(\mathbf{pk}, m) = ([m]_t + \mathbf{r} \cdot \mathbf{u}_0 + \mathbf{t}\mathbf{e}_0]_Q, [\mathbf{r} \cdot \mathbf{u}_1 + \mathbf{t}\mathbf{e}_1]_Q)$;
BFV. $\text{Enc}(\mathbf{pk}, m) = ([\Delta_{\text{BFV}} \cdot [m]_t + \mathbf{r} \cdot \mathbf{u}_0 + \mathbf{e}_0]_Q, [\mathbf{r} \cdot \mathbf{u}_1 + \mathbf{e}_1]_Q)$, where $\Delta_{\text{BFV}} = \lfloor Q/t \rfloor$;
CKKS. $\text{Enc}(\mathbf{pk}, m) = ([m + \mathbf{r} \cdot \mathbf{u}_0 + \mathbf{e}_0]_Q, [\mathbf{r} \cdot \mathbf{u}_1 + \mathbf{e}_1]_Q)$.
- Decryption. Given a secret key $\mathbf{sk} = (1, \mathbf{s})$ and a ciphertext $\mathbf{ct} = (\mathbf{c}_0, \mathbf{c}_1) \in R_{Q'}^2$, the decryption algorithm is like the following:

BGV. $\text{Dec}(\mathbf{sk}, \mathbf{ct}) = [\mathbf{c}_0 + \mathbf{c}_1 \cdot \mathbf{s}]_{Q'}$;
BFV. $\text{Dec}(\mathbf{sk}, \mathbf{ct}) = \lfloor t/Q' \cdot [\mathbf{c}_0 + \mathbf{c}_1 \cdot \mathbf{s}]_{Q'} \rfloor$;
CKKS. $\text{Dec}(\mathbf{sk}, \mathbf{ct}) = [\mathbf{c}_0 + \mathbf{c}_1 \cdot \mathbf{s}]_{Q'}$.

Note that for leveled schemes, the modulus Q' is equal to Q_l at level l , while for the scale-invariant scheme BFV, the modulus will not change.
- Evaluation. Homomorphic encryption provides a way to perform operations on ciphertext without decryption. Common homomorphic evaluations include homomorphic addition and multiplication. For the three schemes, we give their detailed evaluation procedures below.
 - Addition. For any BGV, BFV, or CKKS schemes, given two ciphertexts \mathbf{ct} and \mathbf{ct}' in $R_{Q'}^2$, their sum is defined identically as the follows:

$\text{Add}(\mathbf{ct}, \mathbf{ct}') = [\mathbf{ct} + \mathbf{ct}']_{Q'}$.
 - Multiplication. Given two ciphertexts $\mathbf{ct} = (\mathbf{c}_0, \mathbf{c}_1)$ and $\mathbf{ct}' = (\mathbf{c}'_0, \mathbf{c}'_1)$ in $R_{Q'}^2$, their product is defined as the follows:

$\text{Mult}(\mathbf{ct}, \mathbf{ct}') = [(\tilde{\mathbf{c}}_0, \tilde{\mathbf{c}}_1) + \lfloor p^{-1} \cdot \tilde{\mathbf{c}}_2 \cdot \mathbf{evk} \rfloor]_{Q'}$.
Here $(\tilde{\mathbf{c}}_0, \tilde{\mathbf{c}}_1, \tilde{\mathbf{c}}_2) = [(\mathbf{c}_0 \cdot \mathbf{c}'_0, \mathbf{c}_0 \cdot \mathbf{c}'_1 + \mathbf{c}_1 \cdot \mathbf{c}'_0, \mathbf{c}_1 \cdot \mathbf{c}'_1)]_{Q'}$ for BGV and CKKS, and $[\lfloor t/Q' \cdot (\mathbf{c}_0 \cdot \mathbf{c}'_0, \mathbf{c}_0 \cdot \mathbf{c}'_1 + \mathbf{c}_1 \cdot \mathbf{c}'_0, \mathbf{c}_1 \cdot \mathbf{c}'_1) \rfloor]_{Q'}$ for BFV.

2.3 Basic FHE operations

The basic operations of RLWE-based fully homomorphic encryption schemes performs on polynomials. We provide here the detailed procedures of some of these functions, which enable the efficiency of the schemes.

Barrett reduction To perform a 128-bit integer modular reduction with a 64-bit storage unit, we implement modified Barrett reduction [12]. Given a modulus q , a precomputed value $\gamma = \lfloor 2^{128}/q \rfloor$, and an integer $a < 2^{128}$, a and γ are represented as $a = 2^{64} \cdot a_1 + a_0$ and $\gamma = 2^{64} \cdot \gamma_1 + \gamma_0$, respectively ($a_1, a_0, \gamma_1, \gamma_0$ are all 64-bit integers). Then, the reduction computes

$$a - \lfloor \frac{a_1 \gamma_1 \cdot 2^{128} + (a_1 \gamma_0 + a_0 \gamma_1) \cdot 2^{64} + a_0 \gamma_0}{2^{128}} \rfloor \cdot q,$$

which equals to

$$a - \lfloor \frac{(a_1 \cdot 2^{64} + a_0) \cdot (\gamma_1 \cdot 2^{64} + \gamma_0)}{2^{128}} \rfloor \cdot q = a - \lfloor \frac{a \cdot \gamma}{2^{128}} \rfloor \cdot q.$$

This result is actually either $[a]_q$ or $[a]_q + q$, and then we can get the final reduced result in a constant-time manner via a conditional subtraction by q .

Number theoretic transform To include fast polynomial multiplication, Number Theoretic Transform (NTT) is utilized. In our setting, polynomials are multiplied modulo $x^N + 1$, and thus we perform a negacyclic version of NTT accordingly. To be specific, for a polynomial $\mathbf{f} \in R_q$ with $2N \mid (q-1)$, there exists a primitive $2N$ -th root of unity $\omega \in \mathbb{Z}_q$, and the negacyclic NTT transform of $\mathbf{f} = (f_0, f_1, \dots, f_{N-1})$ is defined as

$$\hat{\mathbf{f}} = (\hat{f}_0, \hat{f}_1, \dots, \hat{f}_{N-1}), \hat{f}_i = \sum_{j=0}^{N-1} f_j \omega^{j(2i+1)}, i \in [0, N).$$

And the inverse, INTT, is performed by replacing the term ω here with ω^{-1} and we omit it for simplicity. Upon the NTT technique, the product of two polynomials $\mathbf{f}, \mathbf{g} \in R_q$ can be computed via $\text{INTT}(\text{NTT}(\mathbf{f}) \circ \text{NTT}(\mathbf{g}))$, where \circ denotes element-wise multiplication of vector.

Residue number system As we mentioned above, the coefficient modulus Q_l is set to be a product of several small pairwise coprime values q_0, \dots, q_l , with each q_i being a prime number itself, so that the Chinese Remainder Theory (CRT) implies a ring isomorphism $R_{Q_l} \xrightarrow{\sim} R_{q_0} \otimes \dots \otimes R_{q_l}$. This means we can replace the arithmetic in R_{Q_l} with several ones with much smaller operands. In practice, there are some RNS instantiations, such as the BEHZ instantiation [11] that proposed to eliminate the multi-precision arithmetic introduced by the division and rounding operation in the decryption and multiplication of BFV.

Modulus switching In homomorphic encryption schemes, the increase in the noise magnitude of a ciphertext accumulates with the increase in homomorphic multiplication operations it goes through. To reduce the noise or to accelerate the homomorphic multiplication computation, the modulus switching procedure is introduced. It performs by scaling or transforming a ciphertext \mathbf{ct} in R_Q into another ciphertext \mathbf{ct}' in R_P , where $Q = \prod_i^l q_i, P = \prod_i^v p_i$ are two coprime moduli. The latter can be efficiently computed by a fast base conversion technique suggested in [11], formulated as

$$\text{Conv}_{Q \rightarrow P}(x \in \mathbb{Z}_Q) = \left(\left[\sum_{i=0}^l \left[x \frac{q_i}{Q} \right]_{q_i} \times \frac{Q}{q_i} \right]_{p_j} \right)_{j=0}^v.$$

For other RNS instantiations, there are similar conversion formulations, and we omit them here for simplicity. For those who are interested, referring to [29], [34] or [31] can help.

2.4 CUDA programming model

The GPU, a multi-core processor with highly parallel features, can execute the same program over multiple data in parallel through the interface Compute Unified Device Architecture (CUDA), which allows access to the GPU resource. In this programming mode, the CPU calls a kernel, which is executed in parallel by several CUDA threads.

CUDA follows a hierarchy, with both execution units and memory. The minimum execution unit is a thread, and multiple threads (up to 1024) can be organized into 1- to 3-dimension, thus forming a block. Similarly, blocks can be arranged into 1- to 3-dimension grid. The purpose of this hierarchy is to be compatible with the structure of operands, making the invocation of elements simpler. GPU contains three types of read-write memory, i.e., global memory (GMEM), shared memory (SMEM) and register file (RF), and two types of read-only memory, i.e., constant memory and texture memory. Threads within a block have access to the block SMEM, and each thread has private RF. The GMEM and read-only memory is accessible to all threads, while the latency is higher contrasted with the SMEM and RF. During the execution of a kernel, each block has to be executed independently in a sequence. Consecutive 32 threads in a block belong to the same warp, which is scheduled by the warp scheduler (WS) in the streaming multiprocessor (SM).

3 Implementation Overview

3.1 Library structure

Our CARM library is proposed to accelerate FHE schemes on both middle-layer and cloud devices, to provide efficient homomorphic evaluation of privacy data. The structure of CARM is shown in Figure 2. Our design follows the host-device (HD) computation model, where the CPU launches a kernel and the GPU computes and returns the result. In functionality, CARM can be divided into two parts, one contains the initialization of the GPU memory and pre-computation of some constants according to the scheme parameters, such as the NTT table; the other contains the implementation of the functions of the FHE scheme, which consists of three layers: the low-level arithmetic layer, polynomial operations layer and the scheme layer.

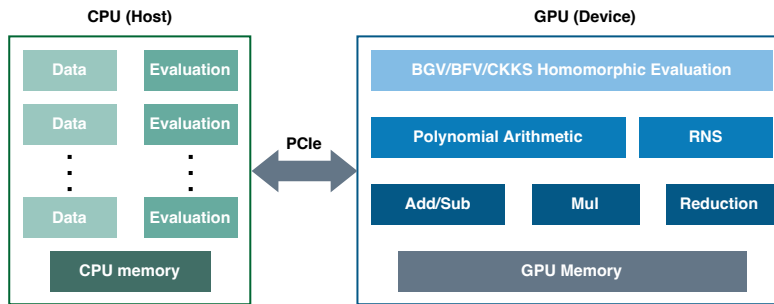


Fig. 2: The host-device model and library structure of CARM.

The basic layer contains the modulo reduction, addition, and multiplication. We provide constant-time implementation written in CUDA parallel thread execution (PTX) instructions and well optimize it by minimizing the instruction numbers and register usage. The middle layer consists of polynomial arithmetic and RNS tools to accelerate the computation of the scheme functions. At the top layer, CARM offers optimized homomorphic evaluation kernels for BGV, BFV and CKKS. Additionally, our library supports $|N| \in [11, 17]$, which covers the parameter requirements of most applications.

3.2 Design considerations and trade-offs

In this section, we illustrate the considerations and trade-offs in the design of CARM. We concentrate on reducing the data access and synchronization and provide many trade-offs between the cost of computation and resource.

Memory-centric optimization The computation of FHE functions is in nature memory-bound since they feature low arithmetic intensity but require plenty of data access. In this work, we exploit memory-centric optimizations to reduce both memory usage and IO latency. For the low-level arithmetic kernel, we make full use of the RF and SMEM and minimize the number of temporary registers. Based on this, we fuse multiple kernels to large ones for high-level homomorphic operations to reduce data transfers, especially the GMEM accesses. Note that after fusing, a function may still contain more than one kernel. While this preserves some data transfer between GMEM and SMEM or RF, it allows a certain degree of parallelism and prevents resource overuse or overflow that lead to low occupancy and performance degradation.

Synchronization Inserting a synchronization point is a commonly applicable method to GPUs to prevent data conflicts when data interaction exists. Throughout this work, we arrange the order of reading and writing data appropriately to eliminate some expensive synchronization. In particular, there is no need to synchronize the 32 threads in the same warp because there is no data conflict. Considering this, we eliminate the thread synchronization in five layers in (I)NTT through finely tune the inner loop.

ILP and TLP For memory-bound computation, reusing data and having one thread execute more operations provides a way to increase the instruction-level parallelism (ILP) and reduce IO latency, thus improving the overall performance. Limited by the fixed on-board storage on the GPU, an SM cannot execute multiple blocks with high memory overhead simultaneously, resulting in some Ws remaining in the idle state thus decrease the thread-level parallelism (TLP). However, low SM occupancy does not always lead to low performance, as storing more data in the SMEM may reduce access to the GMEM, which may consume much more cycles than arithmetic operations. In this work, we choose suitable ILP and TLP for implementations of different parameters, which lead to peak kernel performance. For example, in the (I)NTT, we instantiate the two kernels with different ILP and form the kernels \mathcal{K}_{l_1} and \mathcal{K}_{l_2} as 8 and 2 per-thread implementation respectively. Thus, we can reduce the memory interaction in \mathcal{K}_{l_1} and reserve more RF in \mathcal{K}_{l_2} to store the temp element in the generic RNS multiplication kernel, which contains \mathcal{K}_{l_2} through kernel fusing, as described in Sec 4.4.

4 Implementation Details and Optimizations

In this section, we present the implementation details of CARM, as well as our optimization techniques and trade-offs. First, we give the constant-time low-level arithmetic implementation written in the CUDA PTX instructions, based on which we implement the polynomial operations, including the (I)NTT and the RNS tools. Then, we optimize the homomorphic multiplication using techniques such as kernel fusing. We strive for a generic, scalable and parametric design that is adaptable according to the platforms, and we provide several strategies for making choices.

4.1 Constant-time arithmetic

The efficient implementation of the low-level arithmetic ensures the performance of the schemes, and keeping the execution time constant provides the ability to resist side-channel attacks. Here, we provide our optimized inline functions, including multiply-accumulation, conditional subtraction, and Barrett reduction and show our approaches to minimize the number of instructions and register usage.

Algorithm 1 Multiply_Accumulate

Input: $acc_1, acc_0, a, b \in [0, 2^{64})$ **Output:** $acc'_1, acc'_0 \in [0, 2^{64})$ 1: `mad.lo.cc.u64` acc_0, a, b, acc_0 $\triangleright acc'_0 = acc_0 + (a \cdot b)_{lo}$ 2: `madc.hi.u64` acc_1, a, b, acc_1 $\triangleright acc'_1 = acc_1 + (a \cdot b)_{hi} + carry$

Multiply-accumulation The multiplication of two 64-bit elements requires two PTX instructions on the GPU, so that the two 64-bit halves of the result are obtained and stored separately. For the multiply-accumulate operation, unlike other methods that perform addition after multiplication, we use the `mad` and `madc` instructions to simplify the process, which is given in Algorithm 1. Specifically, we split the 128-bit addition and merge it with the 64-bit multiplication. The other issue that needs to be addressed is the carry. We store the carry flag implicitly in the low 64-bit multiply-accumulate operation, and then use the `madc` instruction to add the carry flag in the high 64-bit multiply-accumulate operation. In this way, we get the result without introducing additional overhead, as the number of instructions used is the same as the original multiplication.

Conditional subtraction Microsoft SEAL set a logic branch to control the condition in subtraction-based reduction. Namely, it reduces numbers from $[0, 2q)$ to $[0, q)$ by subtracting q from elements greater than q and remaining the others unchanged. However, this unbalanced execution poses a potential threat of side-channel attacks and introduces thread divergences in warps, which will incur a drop in the instruction throughput as the different execution paths must be serialized. The instructions with a guard predicate are commonly used for implicit comparisons. Nevertheless, the predicate variable prefixed to the next instruction still specifies a conditional execution. To mitigate the problems, we follow [14] to exploit the arithmetic and shift instructions and subtract q from the input, and the sign gives the implicit comparison result, which can be obtained by logic shifting. Through this approach, we can conduct a conditional subtraction in constant-time and eliminate the divergence in warps caused by the previous control flow instructions.

Barrett reduction The main idea of Barrett reduction is to transform a big integer division to fast arithmetic like multiplication and shifting. The multiplication of two 64-bit elements (i.e., $a = x \cdot y$) yields a 128-bit product, so we implement both 64-bit and 128-bit Barrett reduction using CUDA PTX ISA. In detail, CARM offers three types of Barrett reduction, one for reducing 64-bit numbers and two others for reducing 128-bit numbers, one of which is a faster variant proposed by Shoup [41] to reduce the computation overhead. For a pseudocode description of our implementation, see Algorithm 2. The reduction of a 64-bit input a follows a conventional and commonly applicable way of computing $tmp_0 = \lfloor \frac{a \cdot \gamma_{64}}{2^{64}} \rfloor$ and subtract q multiple of it from a . However, when reducing the product of two 64-bit numbers, we can only compute and store the high 64-bit and low 64-bit of the 128-bit number and *ratio* separately, due to the limitation of CUDA computing capability. Here, we present a constant-time approach by using the combination of instruction `mul.hi.u64` and `mad.lo.cc.u64` to implicitly record the overflow flag in the condition code register and handle the carry bits of additions. The faster version of the 128-bit reduction takes inputs of x and a precomputed number $\lfloor \frac{y \cdot 2^{64}}{q} \rfloor$, which trades the storage of $\lfloor \frac{y \cdot 2^{64}}{q} \rfloor$ for simplifying the computation of $\lfloor \frac{x \cdot y \cdot \gamma_{128}}{2^{128}} \rfloor$. As the output is in range $[0, 2q)$, we use the conditional subtraction approach illustrated before to ensure a constant-time execution. The ratio $\gamma_i = \lfloor \frac{2^i}{q} \rfloor$ is precomputed and stored in the GMEM. During the computation, we load the ratio into the RF and all the calculations are performed in the RF to obtain the minimum latency. Additionally, we reduce the use of temporary registers and reduce the number to one to minimize memory usage.

Algorithm 2 Barrett_Reduction_128

Input: $a = 2^{64} \cdot a_1 + a_0, \gamma = 2^{64} \cdot \gamma_1 + \gamma_0, q, 0 \leq a, \gamma, q < 2^{64}$ **Output:** $res = a \bmod q, 0 \leq res < 2q$

```
1:  mul.hi.u64  tmp, a0, γ0
2:  mad.lo.cc.u64 tmp, a0, γ1, tmp
                                     ▷  $c_0 = \lfloor ((a_0r_0)_{hi} + (a_0r_1)_{lo}) / 2^{64} \rfloor$ 
3:  madc.hi.u64  res, a0, γ1, 0
4:  mad.lo.cc.u64 tmp, a1, γ0, tmp
                                     ▷  $c_1 = \lfloor ((a_0r_0)_{hi} + (a_0r_1)_{lo} + (a_1r_0)_{lo}) / 2^{64} \rfloor$ 
5:  madc.hi.u64  res, a1, γ0, res
6:  mad.lo.u64   res, a1, γ1, res
                                     ▷  $res = (a_0r_1)_{hi} + (a_1r_0)_{hi} + (a_1r_1)_{lo}$ 
7:  mul.lo.u64   res, res, q
8:  sub.u64     res, a0, res
                                     ▷  $res = a_0 - (res * q)_{lo}$ 
9:  sub.s64     res, res, q
10: shr.s64    tmp, res, 63
11: and.b64    tmp, tmp, q
12: add.s64    res, res, tmp
```

4.2 NTT implementation and trade-offs

A full transformation of a $(N - 1)$ -degree polynomial takes $\log N$ layers. Each layer contains the processing of $N/2$ butterfly operations, of which the inputs are two coefficients and the corresponding root. Throughout this paper, we apply the hierarchical implementation [30,33] and divide the procedure into two steps, yielding two kernels that comprise an (I)NTT execution and are responsible for processing different layers respectively, namely, l_1 and l_2 layers, where $\log N = l_1 + l_2$. For compatibility, we provide a generic GPU implementation of (I)NTT that works for all cyclotomic rings with power-of-two orders in the range $\{2^{11}, 2^{12}, \dots, 2^{17}\}$. Additionally, we implement two versions of (I)NTT, one focusing on better performance and the other on using less resource, taking into account the features of different platforms in IoT scenarios.

Previous work shows that setting l_1 and l_2 close to $\frac{1}{2} \log N$ is the most efficient method [10], however, the IO latency is not taken into account. For example, in [33] the authors fixed $l_1 = 8$ and obtain 14- to 17-layer (I)NTT implementation by adjusting l_2 . In each Kernel, they instantiate it with 256 threads and used a fixed 8 per-thread implementation, which requires the execution of a data interaction between RF and SMEM every three layers. Considering that access to RF is more efficient than SMEM, we explore a mix-type memory usage strategy, which is different from the prior implementations [22,33]. Figure 3 shows the process and data access pattern of the two kernels.

Performance-prior version Let λ denote the batch size of RNS-NTT and \mathcal{K}_l denote the kernel that performs l -point (I)NTT. We instantiate the \mathcal{K}_{l_1} and \mathcal{K}_{l_2} kernels with $\lambda N / \rho$ threads, where each thread performs ρ per-thread in-place (I)NTT. The \mathcal{K}_{l_1} kernel is responsible for executing 1 to 6 layers of the NTT. For those that require only three layers or less, we load the data directly from GMEM into RF and let $\rho = 2, 4,$ and 8 corresponding to the cases where $l_1 \in \{1, 2, 3\}$, respectively. For $l_1 \in \{4, 5, 6\}$, we let $\rho = 8$ and exchange the data in the RF through the SMEM. In the \mathcal{K}_{l_2} kernel, we set $l_2 = 11$, store the data in SMEM, and fix $\rho = 2$. With this instantiation, we can set aside more RF resources, making other operations faster and with lower latency.

Memory-prior version To perform an N -point NTT with a batch size β , it requires a $N\beta/128$ KB and $N\beta/64$ KB of storage, respectively, to hold the coefficients and precomputed table of `uint_64` type, e.g., 3β MB in total in the case $N = 2^{17}$. However, there may not be enough resources for storage-constrained devices, making the execution impossible. Several prior works [35,37] proposed generating the twiddle factors in an on-the-fly manner to reduce the size of precomputed table. In detail, instead of generating all twiddle factors,

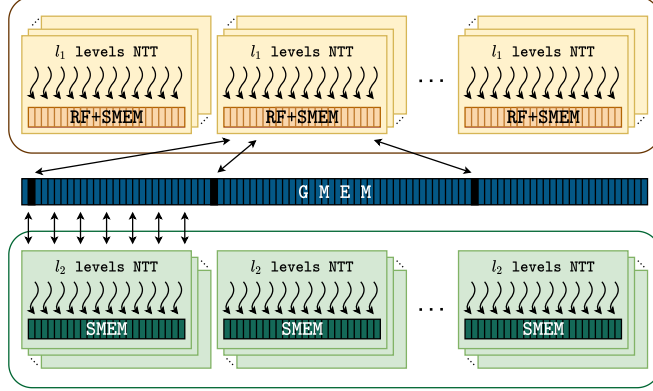


Fig. 3: Execution and memory access pattern of the proposed (I)NTT kernel.

this method only needs to store some of them in advance, and the rest can be generated by these. In [35], the authors applied a base-1024 approach on the last one or two layers, and the number of precomputed twiddle factors can be at most reduced to $(N/4 + (1024 + N/1024))\beta$, in which the computation of $\omega \times x$ is performed by calculating $x' = \omega_1 \times x$ and $x = \omega_2 \times x'$, where $\omega = \omega_1\omega_2$. However, this approach does not minimize the table size.

In our implementation of the memory-prior version, we follow [35] while exploiting a base- \mathcal{D} approach, where \mathcal{D} is a power of 2. We adjust the value of \mathcal{D} to $\log \mathcal{D} = \frac{1}{2} \log N$ so that the number of precomputed twiddle factors is optimal for every N , i.e., $\mathcal{D} + N/\mathcal{D}$. Taking $N = 2^{17}$ as an example, in which case $\mathcal{D} = 256$, through this method, we can reduce the precomputed table size from 2β MB to 12β KB, which means a $170.7\times$ improvement.

Constant-time PTX butterfly For the butterfly operation, we apply the Harvey’s algorithm [32] and build inline device functions with CUDA PTX instruction set. To eliminate the modular multiplication of n^{-1} in the INTT, we adopt the approach following [33,38], which modifies the precomputed table of INTT by multiplying with $\frac{1}{2}$. In GS butterfly this modification only affects one operand, thus previous works [33,38] performed division by 2 on the other operand with a 1-bit shift to the right, where odd numbers need to be added with q to ensure correctness. While this is efficient, it is risky because the parity of the current operand can be inferred through energy analysis, thereby leaking information. To solve this, we multiply the least significant bit of the operand by q and add the result to it. This introduces one multiplication, but ensures a constant-time execution.

4.3 Fast basis conversion

The fast basis conversion $\text{Conv}_{Q \rightarrow P}$ [11] offers an efficient approach to convert the basis of the residues of a polynomial, of which the new basis is coprime to the original basis Q . In the SEAL library, the polynomials are stored in RNS representation, which forms a matrix of size $|Q| \times N$, and the addresses of the residues of the same modulus are contiguous. Let \mathbf{I} denote the input matrix of residues basis Q , \mathbf{O} denote the output matrix of residues under basis P , and \mathbf{C} denote the conversion matrix. The SEAL library conducts an implicit matrix transposition through the iterators and obtains the conversion result through computing $\mathbf{I}^T \cdot \mathbf{C} = \mathbf{O}^T$. However, this transposition triggers the problem of non-unit-stride global memory access on the GPU, which should be avoided because of its high IO latency.

For this consideration, we form a 3-dimension kernel with N threads, where the x-, y-, and z-axis are set based on the block size, $|Q|$, and ciphertext size, i.e., the number of polynomials, respectively. Each thread takes a row of the \mathbf{C} matrix and a column of the \mathbf{I} matrix, performs the inner product by an internal loop, and stores the result in the RF. We use two 64-bit registers to form a 128-bit accumulator to store the result.

Algorithm 3 CT Butterfly in CUDA PTX Assembly

Input: $0 \leq X, Y < 4q$, $\varpi = 2^{64}$, $q < \varpi/4$, $0 < W < q$, $W' = \lfloor W\varpi/q \rfloor$ **Output:** $0 \leq X', Y' < 4q$

```
1:  sub.s64  X, X, 2q
2:  shr.s64  T, X, 63
3:  and.b64  T, T, 2q
4:  add.s64  X, X, T
5:  mul.hi.u64 T, Y, W'
6:  mul.lo.u64 T, T, q
7:  mul.lo.u64 V, Y, W'
8:  sub.u64  V, V, T
9:  add.u64  Y, X, 2q
10: sub.u64  Y, Y, V
11: add.u64  X, X, V
```

$\triangleright (X > 2q) ? X - 2q : X$

$\triangleright V = [WY]_{\varpi} - [\lfloor \frac{W'Y}{\varpi} \rfloor \cdot q]_{\varpi}$

$\triangleright Y' = X - V + 2q$
 $\triangleright X' = X + V$

Algorithm 4 GS Butterfly in CUDA PTX Assembly

Input: $0 \leq X, Y < 2q$, $\varpi = 2^{64}$, $q < \varpi/4$, $0 < W < q$, $W' = \lfloor W\varpi/q \rfloor$ **Output:** $0 \leq X', Y' < 2q$

```
1:  add.u64  T, X, 2q
2:  sub.u64  T, T, Y
3:  add.u64  X, X, Y
4:  sub.s64  X, X, 2q
5:  shr.s64  V, X, 63
6:  and.b64  V, V, 2q
7:  add.s64  X, X, V
8:  and.b64  V, T, 1
9:  mul.lo.u64 V, V, q
10: add.u64  X, X, V
11: shr.s64  X, X, 1
12: mul.hi.u64 V, T, W'
13: mul.lo.u64 V, V, q
14: mul.lo.u64 Y, T, W'
15: sub.u64  Y, Y, V
```

\triangleright Judge parity

$\triangleright X' = \frac{(X+Y)}{2}$

$\triangleright Y' = \frac{W(X-Y)}{2}$

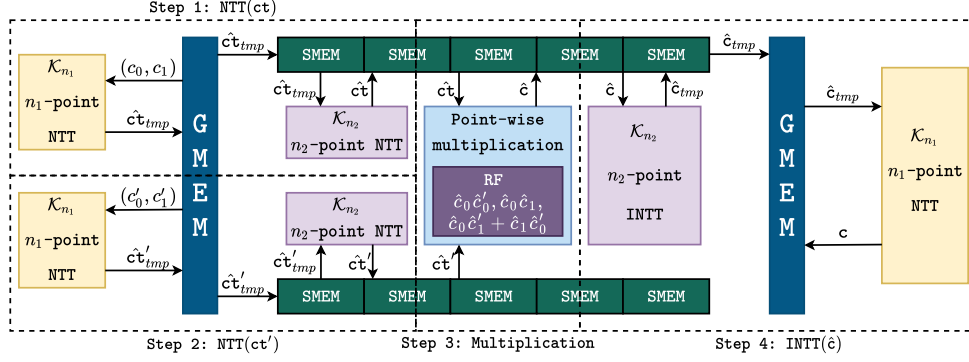


Fig. 4: Fused homomorphic multiplication kernel.

A special case is that in the BEHZ variant of BFV [11], it requires converting the residues from base Q to another base extended by an extra modulus \tilde{m} . In SEAL, it is performed by calling two individual conversions respectively. We combine these two processes in our GPU implementation, i.e., a thread takes the residues modulo p_i and \tilde{m} , and one column of the \mathbf{I} matrix, performs two vector inner products and stores the results in two 128-bit accumulators, respectively. This approach reduces the memory access but increases the computational overhead of the threads, since the product of the residues of modulo \tilde{m} by the \mathbf{I} matrix is repeated by all threads, but due to the parallelism of the computation, it introduces no additional time consumption.

4.4 Homomorphic multiplication

In the SEAL library, polynomials are represented in different forms, i.e., only CKKS keeps the elements in the NTT form. To reduce the computation overhead, the (I)NTT should be conducted for the ciphertexts of BGV and BFV. Meanwhile, the design of BFV makes it require more operations in multiplication, such as basis conversion and scale, which is different from BGV. For compatibility, we exploit a generic RNS multiplication kernel (GRM) that works for both BGV and BFV, which fuse the NTT, element-wise product and INTT kernels into one kernel so that the interaction with GMEM can be reduced to the load and store operations required by the (I)NTT themselves. The logic flow and data access pattern is shown in Figure 4.

Generic RNS multiplication kernel Let $\hat{ct} = \text{NTT}(ct)$ and the variables containing subscript tmp be the intermediate value in the (I)NTT processing. In our GRM kernel, first, we load the ciphertexts $ct = (c_0, c_1)$ and $ct' = (c'_0, c'_1)$ from GMEM and transform them into NTT domain using the \mathcal{K}_{n_1} and \mathcal{K}_{n_2} kernel defined before. Then, each thread loads data to the RF, computes the triple $\hat{c} = ([c_0 \cdot c'_0]_{Q_1}, [c_0 \cdot c'_1 + c_1 \cdot c'_0]_{Q_1}, [c_1 \cdot c'_1]_{Q_1})$ and stores the results back to the SMEM. After that, we perform INTT to the elements in the triple respectively. During the entire process, the remaining load and store operations over GMEM are that required in performing (I)NTT on ct , ct' and c , which in nature cannot be omitted. To fully utilize the RF, we instantiate the \mathcal{K}_{n_2} kernel with $n_2 = 11$ that forms the 2 per-thread in-place n_2 -point NTT. Through this configuration we can reserve enough registers for the computation of the triple, which ensures no register overflow will occur that impact the correctness in launching kernel.

One thing to note is that, for implementations with different parameters, especially N , kernel fusing does not always improve the performance. In the original procedure, i.e., multiplication is composed by several kernels containing the NTT, element-wise product and INTT, where blocks are scheduled by different SMs and executed in parallel. However, in the GRM kernel, they are executed sequentially. Therefore, for high performance platforms with a large number of SMs and a small N , it is more efficient not to use kernel fusing. In order to obtain the bounds between the two methods, we have performed adequate experiments and given them in detail in Sec 5.3.

Table 1: Testing environment.

Target platform	Middle-layer device	Cloud device
GPU	Jetson AGX Xavier	Tesla V100S PCIe
GPU cores	512	5120
GPU frequency	1.37 GHz	1.60 GHz
SMs	8	80
GPU memory	32 GB (unified)	32 GB
Compute capability	7.2	7.0
CUDA version	11.4	11.6
CPU	NVIDIA Carmel	Intel Xeon Silver
	64-bit ARMv8.2	4210R
CPU cores	8	10
CPU frequency	2.26 GHz	2.40 GHz
CPU memory	32 GB (unified)	32 GB
CPU Cache	8 MB L2	10 MB L2
	4 MB L3	13.75 MB L3

Kernel fusing in BFV Our generic design is also applicable to the full RNS homomorphic multiplication applied in the BEHZ variant of BFV [11], which is implemented in SEAL. Here, we follow the notation in [11], and denote the extended base as $\mathcal{B} \cup \{m_{\text{sk}}\} \cup \{\tilde{m}\}$, or $\mathcal{B}_{\text{sk}} \cup \{\tilde{m}\}$, where $\mathcal{B}_{\text{sk}} = \mathcal{B} \cup \{m_{\text{sk}}\}$, \mathcal{B} is the auxiliary base with v moduli, and m_{sk} and \tilde{m} are two extra moduli. The main function of this process (described in Algorithm 3, [11]) is the fast base conversion, which is performed to compute the product in an extended basis and to achieve the approximate rounding and transitional step. Considering the structure of the operands and memory usage, we exploit the kernel fusing (KF) approach and fuse the entire process into four kernels. Note that kernel fusing allows data reuse that saves data access instructions. Fusing a function consists of more than one kernel, will reserve enough memory resources for each kernel, thus reducing the overall IO latency.

In our implementation, first, we fuse the Small Montgomery Reduction into the fast base conversion kernel, using the method described in Sec 4.3. After fusing, the kernel performs the conversion of ct and ct' from base Q_l to $\mathcal{B}_{\text{sk}} \cup \{\tilde{m}\}$, yielding new ciphertexts with q -overflows, which are then reduced for elements in \mathcal{B}_{sk} through the Small Montgomery Reduction. Second, we apply our generic multiplication kernel to multiply the two ciphertexts in the NTT domain, by setting the modulus size to $|Q_l \cup \mathcal{B}_{\text{sk}}|$. Third, we fuse the scaling into the conversion to \mathcal{B}_{sk} . In the last kernel, which performs the Shenoy and Kumaresan like conversion [40], we apply a similar method as the first one and let each thread performs the vector inner products as well as the computation of α_{sk} . In the four kernels, a block consists of 128 threads and loads 128 columns of the \mathbf{I} matrix, which is a multiple of 32, so that all memory write requests of a warp fall into distinct memory banks and will not cause bank conflict.

5 Results and Comparison

5.1 Testing environment

To evaluate the performance of CARM, we implement it on two NVIDIA GPUs, one is the Jetson AGX Xavier developer kit and the other is the Tesla V100S PCIe, targeting the embedded and server-grade platforms. The detailed testing environment hardware configurations and the corresponding platforms are summarized in Table 1. We integrate our GPU implementation into Microsoft SEAL library v4.0 [39] and all implementations are profiled with the same build and execute environment. For all experiments, we compile the C++ implementation using g++ 11.1 and the GPU implementation using CUDA 11 on Ubuntu 20.04.

All the parameters we use all reach 128-bit security according to the LWE estimator [5], except benching the batched (I)NTT with $|N| = 14$ and batch size $\beta = 21$, because we wanted to be consistent with the parameters selected in [35] for better comparisons of performance.

Table 2: The performance of our (I)NTT implementation on Tesla V100S, CPU baseline, and [38].

$ N $	Batch size	Our work (PP)		CPU		[38]	NTT speedup	
		NTT	INTT	NTT	INTT	NTT	vs. CPU	vs. [38]
11	1	12.3	12.3	35.1	24.8	12.5	2.9×	1.0×
12	1	14.4	15.4	71.8	54.0	22.5	5.0×	1.6×
13	1	15.4	16.4	151.2	115.0	27.0	9.8×	1.8×
14	1	17.4	18.4	294.9	242.8	29.0	16.9×	1.7×
15	1	16.4	16.4	633.6	486.1	39.0	38.6×	2.4×
16	1	16.4	17.4	1299.5	1022.7	-	79.2×	-
17	1	19.5	19.5	2746.2	2176.0	-	140.8×	-

Table 3: The performance of our batched (I)NTT implementation on Tesla V100S, CPU baseline, and [35].

$ N $	Batch size	Our work (PP)		CPU		[35]	NTT speedup	
		NTT	INTT	NTT	INTT	NTT	vs. CPU	vs. [35]
12	2	15.3	15.4	142.8	108.0	-	9.3×	-
13	4	15.4	16.4	563.0	431.0	-	36.7×	-
14	21	46.0	49.2	6264.0	4814.3	44.1	136.0×	-
15	21	68.6	71.7	13137.5	10258.9	84.2	191.6×	1.2×
16	21	123.9	132.1	27682.2	21805.0	156.3	223.5×	1.3×
17	21	231.4	240.6	58182.5	46879.1	304.2	251.4×	1.3×

5.2 NTT

Our generic and scalable design of (I)NTT kernels provides several ways of combining \mathcal{K}_{l_1} and \mathcal{K}_{l_2} for different parameters. The $l_2 = 11$ with $\rho = 2$ configuration targets for less SMEM usage with a high SM occupancy, and reserve sufficient RF to store the temp element after fusing multiple kernels into one. Table 2 and 3 present the performance of our implementation of (I)NTT with the performance-prior (PP) configuration on Tesla V100S, which are different in the batch size, and the comparisons with the CPU baseline implemented in SEAL v4.0 [39] and related works [35, 38]. Additionally, we provide the benchmark of both performance-prior (PP) and memory-prior (MP) configuration of (I)NTT on Jetson AGX Xavier, as well as the CPU baseline in Table 4. Since there is no (I)NTT implementation on embedded GPUs yet, we only list the comparison with the CPU implementation. Here, [38] is a recent published work that reports the speed of single (I)NTT with $|N| \in [11, 15]$ on Tesla V100, and in [35] the authors give the execution time of NTT with $|N| \in [14, 17]$ batch size $\beta = 21$ on NVIDIA Titan V, in which the twiddle factors of the last layer are computed on-the-fly.

In the case of performing a single transformation, our GPU implementation provides $2.9\times$ to $254.1\times$ speedup on V100S and $1.1\times$ to $5.7\times$ speedup on AGX Xavier compared to the CPU baseline. Meanwhile, the speedup factor grows rapidly with the polynomial degree and batch size, and we improve the performance up to $140.8\times$ and $254.1\times$ correspondingly for single and batched NTT when $|N| = 17$. Our implementation outperforms [38] and [35] by $1.0\times$ to $2.4\times$, and offers a larger parameter set. Although our implementation is a bit slower than [35] when $(|N|, \beta) = (14, 21)$, this set of parameters does not satisfy the security requirement. Meanwhile, they employ non-constant-time implementation that may leak the parity of operands [33, 35]. Compared to their approach, we offer constant-time implementation of (I)NTT supports $|N| \in [11, 17]$, which is in nature generic and scalable and suitable for many scenarios.

To measure the impact of our MP approach on the size of the precomputed TW tables, we present the results of different implementations in Table 5. Note that this is the size of the table needed for a single NTT, and the results of [35] are given as the smallest size that their implementation can lead to. Our implementation has a significant effect on reducing the table size, by a factor of 21.3 to 170.7 compared to the method without MP. Additionally, we achieve $20.1\times$ to $44.2\times$ improvements against [35].

Table 4: The performance of our batched (I)NTT implementation on Jetson AGX Xavier and CPU baseline.

$ N $	Batch size	Our work (PP)		Our work (MP)		CPU		Speedup
		NTT	INTT	NTT	INTT	NTT	INTT	
11	1	111.9	110.8	150.8	162.0	199.6	122.7	1.8×/1.1×
12	1	153.2	154.8	186.2	210.0	211.0	234.5	1.4×/1.5×
13	1	159.0	170.9	217.2	246.4	372.2	497.8	2.3×/2.9×
14	1	215.8	248.0	282.3	317.3	738.8	856.7	3.4×/3.5×
15	1	308.7	336.2	404.1	459.5	1384.3	1408.4	4.5×/4.2×
16	1	564.9	624.1	769.3	902.7	3011.6	3000.4	5.3×/4.8×
17	1	1079.4	1170.4	1559.8	1801.2	5931.1	6253.2	5.5×/5.3×
12	2	160.2	163.7	196.1	222.7	275.1	272.1	1.7×/1.7×
13	4	375.0	431.3	451.4	519.4	1284.4	1278.5	3.4×/3.0×
14	8	1440.5	1603.9	1801.5	2040.7	6520.8	5512.9	4.5×/3.4×
15	16	3944.3	4070.1	5562.1	6305.6	21108.4	23246.5	5.4×/5.7×

Table 5: The TW table size of the baseline implementation, [35], and our MP method (measured in KB).

$ N $	Baseline	[35]	Our work (MP)	Improvement	
				vs. CPU	vs. [35]
11	32		1.5	21.3×	
12	64	-	2	32.0×	-
13	128		3	42.7×	
14	256	80.25	4	64.0×	20.1×
15	512	144.5	6	85.3×	24.1×
16	1024	273	8	128.0×	34.1×
17	2048	530	12	170.7×	44.2×

5.3 Homomorphic multiplication

In this section, we analyze the performance impact of our implementations using different optimization strategies, as well as the comparisons with related works.

Because we believe that kernel fusing does not improve the performance in all cases, we test the execution time of the two methods to obtain the bound, i.e., with or without GRM, respectively, where the results of $|N| = 14$ and 15 on the Tesla V100S are shown in Fig. 5a and 5b. From our experiments, we find that the method without GRM is faster when $|N|$ is less than 14. This result starts to change at $|N| = 14$ and the method with GRM outperforms the other when the number of moduli is greater than 6. Additionally, GRM shows more advantages when $|N|$ and the modulus count becomes larger. Throughout this work, we apply this bound and combine both methods to get better performance.

Based on this, we evaluate the performance of our optimized homomorphic multiplication of BGV, BFV and CKKS. Table 6 summarizes the execution time on a single-thread CPU, our GPU implementation on Tesla V100S, and some related works, where [4] and [33] are the state-of-the-art implementations of BFV and CKKS, and [6] and [8] are some recent works. The target platforms in these works have computing capability similar to ours. Although these works do not cover as many parameters of fully homomorphic encryption schemes in practice as we do, we still find adequate data for comparison. Additionally, we provide the performance on Jetson AGX Xavier in Table 7. Compared to the CPU baseline, our GPU implementation provides up to 378.4× speedup for BGV, 234.5× speedup for BFV, and 287.2× speedup for CKKS on Tesla V100S, and 8.8×, 9.2×, and 10.3× on Jetson AGX Xavier. Meanwhile, our implementation outperforms the works listed in Table 6, which shows the impact of our optimizations.

5.4 Discussions

Resistance to side-channel leakage The butterfly and reduction operations play important roles in the implementation of HE schemes. Unfortunately, the majority open-source libraries and previous works use

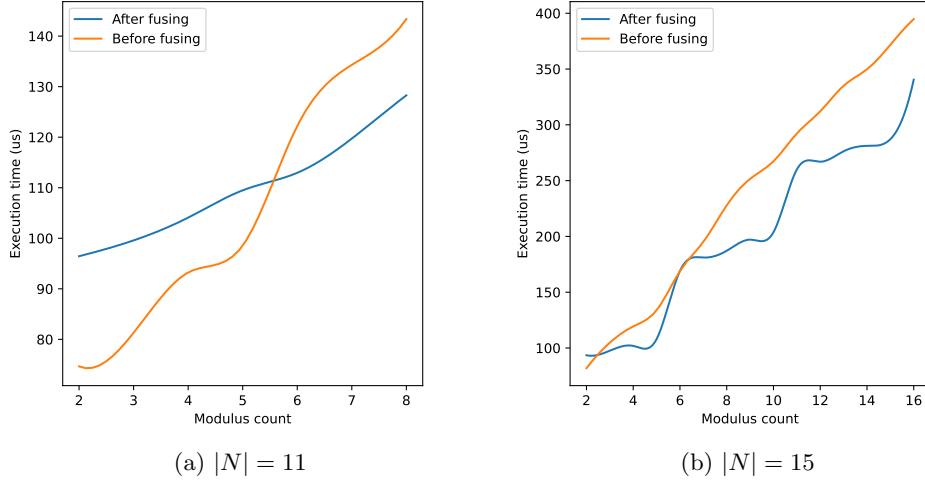


Fig. 5: Performance of the generic multiplication before and after fusing.

Table 6: Performance of homomorphic multiplication on Tesla V100S and comparisons with other works.

$ N $	l	BGV			BFV			CKKS			[4]	[6]	[8]	[33]
		GPU	CPU	Speedup	GPU	CPU	Speedup	GPU	CPU	Speedup	P100	K80	V100	V100
12	2	0.083	1.754	21.1×	0.167	4.049	24.2×	0.024	0.144	6.0×	1.833	1.214	-	
13	4	0.090	5.130	57.0×	0.189	16.213	85.8×	0.023	0.639	27.8×	3.538	3.061	0.403	
14	8	0.127	20.625	162.4×	0.386	71.727	185.8×	0.031	2.487	80.2×	11.747	13.914	0.742	-
15	16	0.355	86.342	243.2×	1.608	365.014	227.0×	0.059	9.988	169.3×			2.388	
16	32	1.080	361.504	334.7×	9.209	2159.234	234.5×	0.157	40.362	257.1×	-	-	33.577	17.4
17	32	1.988	752.253	378.4×	19.463	4546.390	233.6×	0.281	80.710	287.2×				7.96

Table 7: Performance of homomorphic multiplication on Jetson AGX Xavier.

$ N $	l	BGV			BFV			CKKS		
		GPU	CPU	Speedup	GPU	CPU	Speedup	GPU	CPU	Speedup
12	2	1.048	4.363	4.2×	2.880	16.176	5.6×	0.173	0.723	4.2×
13	4	2.502	14.115	5.6×	8.800	65.036	7.4×	0.395	2.186	5.5×
14	8	9.549	56.164	5.9×	37.756	258.971	6.9×	1.006	8.862	8.8×
15	16	24.946	218.688	8.8×	112.884	1043.381	9.2×	3.380	34.726	10.3×

compiler-level optimization as a compensation to eliminate time variance. This poses a potential threat, as the study has shown that simple modifications to the code under this protection can reduce the hardness of the R-LWE problem [24]. In this paper, we use a constant-time design for all low-level arithmetic operations with the least possible performance loss. This technique brings an additional benefit, as an unbalanced execution will introduce thread divergences in warps, making different execution paths have to be serialized. We believe that such a countermeasure is worthwhile in two ways: first, it reduces the potential leakage of secret information; second, it increases the instruction throughput, which contributes to the efficiency.

Applicability of proposed approaches Our design is generic and provides an optimization approach for GPU implementations of other homomorphic schemes and on other platforms. The difference of GPUs is mainly in terms of resource and computational capability. For different fully homomorphic encryption schemes, especially those based on RLWE, the low-level operations have low arithmetic density and share similar constructions. In our implementation, firstly, we propose some general designs with adjustable parameters that are scalable to be chosen according to the specific situation. Second, we provide both performance-prior and memory-prior approaches that can be applied to devices with different computational capabilities. Additionally, we use a memory-centric optimization approach based on kernel fusing that reduces memory accesses. We also provide the impact of kernel fusing on performance, giving a choice between different implementation approaches. In summary, our approach is meaningful for the optimization of schemes on various platforms.

6 Conclusion

In this work, we present CARM, the first optimized GPU implementation of word-wise homomorphic encryption schemes in IoT scenario, covering BGV, BFV and CKKS. Our work focuses on homomorphic multiplication, and offers various trade-offs between computational efficiency and memory consumption for deployment on different platforms. We provide evaluations on two different GPUs, targeting the embedded and cloud devices, and the results demonstrate the effectiveness of our approaches. Our generic and parametric design is applicable to GPU implementations of other functions, schemes, and on other platforms, which will be our future work.

References

1. Heaan. <https://github.com/snucrypto/HEAAN> (Apr 2022)
2. Helib. <https://github.com/homenc/HElib> (Apr 2022)
3. Palisade. <https://gitlab.com/palisade> (Apr 2022)
4. Al Badawi, A., Veeravalli, B., Mun, C.F., Aung, K.M.M.: High-performance fv somewhat homomorphic encryption on gpus: An implementation using cuda. *IACR Transactions on Cryptographic Hardware and Embedded Systems* **2018**(2), 70–95 (2018). <https://doi.org/10.13154/tches.v2018.i2.70-95>
5. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. *J. Math. Cryptol.* **9**(3), 169–203 (2015), <http://www.degruyter.com/view/j/jmc.2015.9.issue-3/jmc-2015-0016/jmc-2015-0016.xml>
6. Alves, P.G.M.R., Ortiz, J.N., Aranha, D.F.: Faster homomorphic encryption over gpgpus via hierarchical dgt. In: *Financial Cryptography and Data Security*. pp. 520–540. *Lecture Notes in Computer Science* (2021). https://doi.org/10.1007/978-3-662-64331-0_27
7. Aydin, F., Karabulut, E., Potluri, S., Alkim, E., Aysu, A.: Reveal: Single-trace side-channel leakage of the seal homomorphic encryption library. *IACR Cryptol. ePrint Arch.* **2022**, 204 (2022)
8. Badawi, A.A., Hoang, L., Mun, C.F., Laine, K., Aung, K.M.M.: Privft: Private and fast text classification with homomorphic encryption. *IEEE Access* **8**, 226544–226556 (2020). <https://doi.org/10.1109/access.2020.3045465>
9. Badawi, A.A., Veeravalli, B., Mi Aung, K.M.: Faster number theoretic transform on graphics processors for ring learning with errors based cryptography. In: *2018 IEEE International Conference on Service Operations and Logistics, and Informatics (SOLI)*. pp. 26–31 (2018). <https://doi.org/10.1109/soli.2018.8476725>
10. Bailey, D.H.: Ffts in external of hierarchical memory. In: *Proceedings of the 1989 ACM/IEEE conference on Supercomputing - Supercomputing '89*. pp. 234–242 (1989). <https://doi.org/10.1145/76263.76288>

11. Bajard, J., Eynard, J., Hasan, M.A., Zucca, V.: A full rns variant of fv like somewhat homomorphic encryption schemes. In: *Selected Areas in Cryptography - SAC 2016*. pp. 423–442. *Lecture Notes in Computer Science* (2017). https://doi.org/10.1007/978-3-319-69453-5_23
12. Barrett, P.: Implementing the rivest shamir and adleman public key encryption algorithm on a standard digital signal processor. In: *Advances in Cryptology - CRYPTO 1986*. *Lecture Notes in Computer Science*, vol. 263, pp. 311–323 (1986). https://doi.org/10.1007/3-540-47721-7_24
13. Boemer, F., Kim, S., Seifu, G., Souza, F.D.M.d., Gopal, V.: Intel hexl: Accelerating homomorphic encryption with intel avx512-ifma52. In: *WAHC '21: Proceedings of the 9th on Workshop on Encrypted Computing & Applied Homomorphic Cryptography*. pp. 57–62 (2021). <https://doi.org/10.1145/3474366.3486926>
14. Bos, J.W., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehlé, D.: Crystals - kyber: A cca-secure module-lattice-based kem. In: *2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018*. pp. 353–367 (2018). <https://doi.org/10.1109/EuroSP.2018.00032>
15. Boura, C., Gama, N., Georgieva, M., Jetchev, D.: Chimera: Combining ring-lwe-based fully homomorphic encryption schemes. *Journal of Mathematical Cryptology* **14**(1), 316–338 (2020). <https://doi.org/10.1515/jmc-2019-0026>
16. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (leveled) fully homomorphic encryption without bootstrapping. In: *Innovations in Theoretical Computer Science 2012*. pp. 309–325 (2012). <https://doi.org/10.1145/2090236.2090262>
17. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) lwe. In: *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011*. pp. 97–106 (2011). <https://doi.org/10.1109/focs.2011.12>
18. Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-lwe and security for key dependent messages. In: *Advances in Cryptology - CRYPTO 2011*. pp. 505–524. *Lecture Notes in Computer Science* (2011). https://doi.org/10.1007/978-3-642-22792-9_29
19. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: *Advances in Cryptology - ASIACRYPT 2017*. *Lecture Notes in Computer Science*, vol. 10624, pp. 409–437 (2017). https://doi.org/10.1007/978-3-319-70694-8_15
20. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Tfhe: Fast fully homomorphic encryption over the torus. *Journal of Cryptology* **33**(1), 34–91 (2020). <https://doi.org/10.1007/s00145-019-09319-x>
21. Cooley, J.W., Tukey, J.W.: An algorithm for the machine calculation of complex fourier series. *Mathematics of computation* **19**(90), 297–301 (1965)
22. Dai, W., Sunar, B.: cuhe: A homomorphic encryption accelerator library. In: *Cryptography and Information Security in the Balkans - Second International Conference, BalkanCryptSec 2015*. *Lecture Notes in Computer Science*, vol. 9540, pp. 169–186 (2015). https://doi.org/10.1007/978-3-319-29172-7_11
23. van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully homomorphic encryption over the integers. In: *Advances in Cryptology - EUROCRYPT 2010*. *Lecture Notes in Computer Science*, vol. 6110, pp. 24–43 (2010). https://doi.org/10.1007/978-3-642-13190-5_2
24. Drucker, N., Pelleg, T.: Timing leakage analysis of non-constant-time ntt implementations with harvey butterflies. *IACR Cryptol. ePrint Arch.* **2022**, 94 (2022)
25. Ducas, L., Micciancio, D.: Fhew: Bootstrapping homomorphic encryption in less than a second. In: *Advances in Cryptology - EUROCRYPT 2015*. *Lecture Notes in Computer Science*, vol. 9056, pp. 617–640 (2015). https://doi.org/10.1007/978-3-662-46800-5_24
26. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. *IACR Cryptol. ePrint Arch.* p. 144 (2012)
27. Gentry, C.: A Fully Homomorphic Encryption Scheme. Ph.D. thesis, Stanford University (2009)
28. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009* (2009). <https://doi.org/10.1145/1536414.1536440>
29. Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the aes circuit. In: *Advances in Cryptology - CRYPTO 2012*. *Lecture Notes in Computer Science*, vol. 7417, pp. 850–867 (2012). https://doi.org/10.1007/978-3-642-32009-5_49
30. Govindaraju, N.K., Lloyd, B., Dotsenko, Y., Smith, B., Manferdelli, J.: High performance discrete fourier transforms on graphics processors. In: *2008 SC - International Conference for High Performance Computing, Networking, Storage and Analysis*. pp. 1–12 (2008). <https://doi.org/10.1109/sc.2008.5213922>
31. Han, K., Ki, D.: Better bootstrapping for approximate homomorphic encryption. In: *Topics in Cryptology - CT-RSA 2020*. pp. 364–390. *Lecture Notes in Computer Science* (2020). https://doi.org/10.1007/978-3-030-40186-3_16
32. Harvey, D.: Faster arithmetic for number-theoretic transforms. *Journal of Symbolic Computation* **60**, 113–119 (2014). <https://doi.org/10.1016/j.jsc.2013.09.002>
33. Jung, W., Kim, S., Ahn, J.H., Cheon, J.H., Lee, Y.: Over 100x faster bootstrapping in fully homomorphic encryption through memory-centric optimization with gpus. *IACR Transactions on Cryptographic Hardware and Embedded Systems* **2021**(4), 114–148 (2021). <https://doi.org/10.46586/tches.v2021.i4.114-148>

34. Kim, A., Polyakov, Y., Zucca, V.: Revisiting homomorphic encryption schemes for finite fields. In: *Advances in Cryptology - ASIACRYPT 2021*. pp. 608–639. *Lecture Notes in Computer Science* (2021). https://doi.org/10.1007/978-3-030-92078-4_21
35. Kim, S., Jung, W., Park, J., Ahn, J.H.: Accelerating number theoretic transformations for bootstrappable homomorphic encryption on gpus. In: *IEEE International Symposium on Workload Characterization, IISWC 2020*. pp. 264–275 (2020). <https://doi.org/10.1109/IISWC50251.2020.00033>
36. López-Alt, A., Tromer, E., Vaikuntanathan, V.: On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In: *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012*. pp. 1219–1234 (2012). <https://doi.org/10.1145/2213977.2214086>
37. Natarajan, D., Dai, W.: Seal-embedded: A homomorphic encryption library for the internet of things. *IACR Transactions on Cryptographic Hardware and Embedded Systems* pp. 756–779 (2021). <https://doi.org/10.46586/tches.v2021.i3.756-779>
38. Özerk, Ö., Elgezen, C., Mert, A.C., Öztürk, E., Savas, E.: Efficient number theoretic transform implementation on gpu for homomorphic encryption. *The Journal of Supercomputing* (2021). <https://doi.org/10.1007/s11227-021-03980-5>
39. Microsoft SEAL (release 4.0). <https://github.com/Microsoft/SEAL> (Apr 2022), microsoft Research, Redmond, WA.
40. Shenoy, A.P., Kumaresan, R.: Fast base extension using a redundant modulus in RNS. *IEEE Trans. Computers* **38**(2), 292–297 (1989). <https://doi.org/10.1109/12.16508>, <https://doi.org/10.1109/12.16508>
41. Shoup, V.: Ntl: A library for doing number theory. Tech. rep. (2001)
42. Smart, N.P., Vercauteren, F.: Fully homomorphic simd operations. *Designs, Codes and Cryptography* **71**(1), 57–81 (2014). <https://doi.org/10.1007/s10623-012-9720-4>